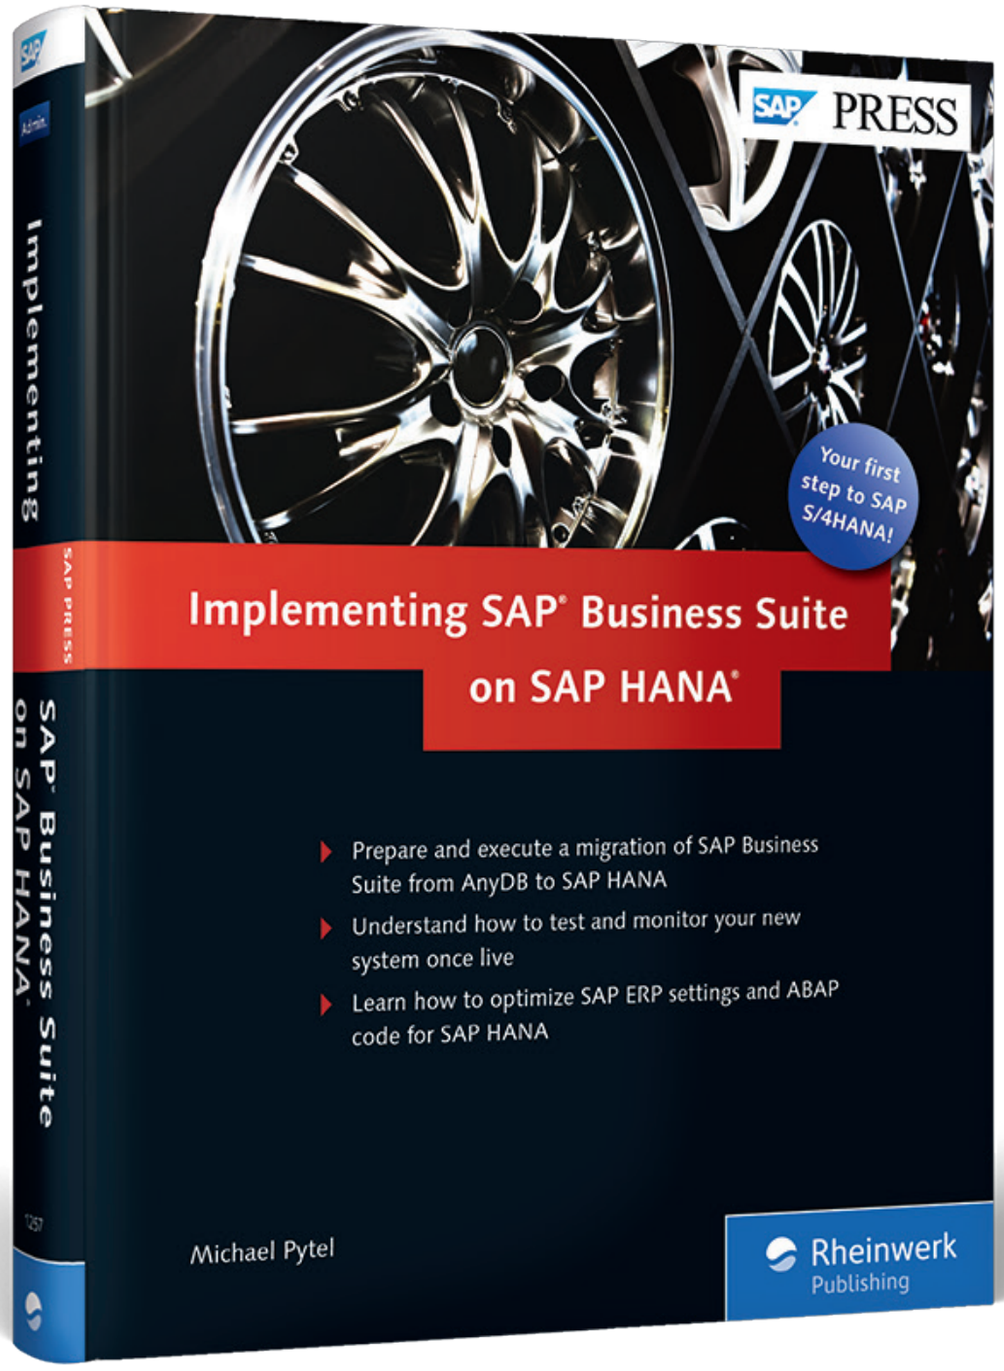


# *Reading Sample*

*This sample chapter describes the infrastructure requirements to prepare your system landscape for the SAP Business Suite on SAP HANA migration. It covers multiple methods for ensuring your environment is sized properly, and demonstrates some of the backup and recovery options available to you.*



Michael Pytel

**Implementing SAP Business Suite on SAP HANA**

597 Pages, 2016, $79.95/€79.95
ISBN 978-1-4932-1257-6

**www.sap-press.com/3895**

*Running SAP Business Suite on SAP HANA requires infrastructure changes to your SAP landscape because new software components and hardware are installed within your landscape. In addition to ensuring that it's sized appropriately, you need to ensure that it's all backed up and recoverable.*

# 3 Infrastructure Planning

Now that we've defined the project scope, you can focus on the decisions you'll need to make about your infrastructure planning. For those customers whose SAP systems began as SAP NetWeaver 7.31 or above, minimal changes are required. For those customers whose SAP ERP systems were initially installed with SAP R/3 or SAP ERP 6.0 on SAP NetWeaver 7.0 and then upgraded over the years, you'll have some more decisions to make. You'll learn more about why SAP is recommending changes to the primary application server (PAS) and how to comply with those recommendations. What is an enqueue server, and why does SAP want it separate? We'll cover this topic and talk about how it affects sizing. Additionally, if you were a customer who installed a Java stack in addition to the ABAP stack during your initial deployment of SAP ERP 6.0, we'll highlight some of the changes required in your landscape before the migration to SAP Business Suite on SAP HANA. SAP doesn't support the migration of a dual stack system to SAP HANA. We'll cover the how and the why as it relates to your SAP Business Suite on SAP HANA migration.

Sizing your SAP Business Suite on SAP HANA environment will depend on the scope you identified in Chapter 2 and the architecture your team plans to deploy. In Chapter 1, you learned about some of the architecture options when running SAP Business Suite on SAP HANA, which we'll expand upon in this chapter by explaining how each architecture option impacts your sizing requirement. SAP supports running the ABAP PAS on the same hardware as the SAP HANA database, which was previously not an option. We'll break down the additive sizing requirements and discuss the potential benefits of this landscape to those customers with small- to medium-sized environments.

In addition to sizing, we'll also explore some of the backup and recovery options available to SAP customers. There are numerous hardware vendor-provided solutions in addition to the capabilities provided by SAP. We'll show you how to find those certified solutions and provide more details on out-of-the-box solutions. This chapter concludes with information on backup and recovery to support the best decision for your organization's landscape.

## 3.1 Enqueue Server

The enqueue server is the function within SAP NetWeaver that maintains data about user locks in the lock table in memory; it keeps users operating day to day in SAP ERP without interrupting other users. When a user opens a purchase order, and the message is displayed that another user has the document locked, this is one function the enqueue server supports. Beginning with SAP NetWeaver 7.31 and continuing in SAP NetWeaver 7.4, SAP began splitting the enqueue services from the central instance by default. This means that new installations you delivered had two instances on the PAS host. One instance supports the ABAP dispatcher, work processes, gateway, Internet Communication Manager (ICM), and Internet Graphics Service (IGS), and the other instance supports the message server and enqueue server. If your system was originally an SAP R/3 system, or the initial install was done using a release prior to SAP NetWeaver 7.31, then all these functions were part of a single instance. There are, of course, exceptions. If your SAP ERP landscape was configured for High Availability (HA), it's likely you already have a standalone enqueue server that is replicated to another host for failover purposes. If you're already using the HA scenario with SAP NetWeaver 7.x, then no action is needed with the enqueue server. If you're a customer with the integrated enqueue server on any version of SAP NetWeaver, then SAP recommends that you split the function from the central instance before the migration to SAP HANA. SAP specifically says:

> *SAP strongly recommends using the standalone enqueue server as the better solution. The standalone enqueue server offers better performance and better scalability for large systems (two or more dialog instances). It is already the standard for all new installations and mandatory for high availability setups. In future releases, it will become mandatory for all systems.*





**Note** [«]

If you want to read SAP's official statement in full, see SAP Note 2019532 – Performance of Integrated Enqueue Server, and SAP Note 2013043 – Performance Problems with Enqueue Work Process.

SAP does provide the tools you need to initially split the enqueue server from the central instance. Using the Software Provisioning Manager 1.0 (SWPM), you can split the instance with minimal application server downtime. Ideally, this is done in the weeks or months before the SAP Business Suite on SAP HANA migration. It's a relatively low-risk activity because you're not moving services to another server. You're keeping the services local but splitting them from the primary instance for performance reasons. In Figure 3.1, you can visualize the changes being made to your system. The update also requires edits to the instance profile for the application servers.

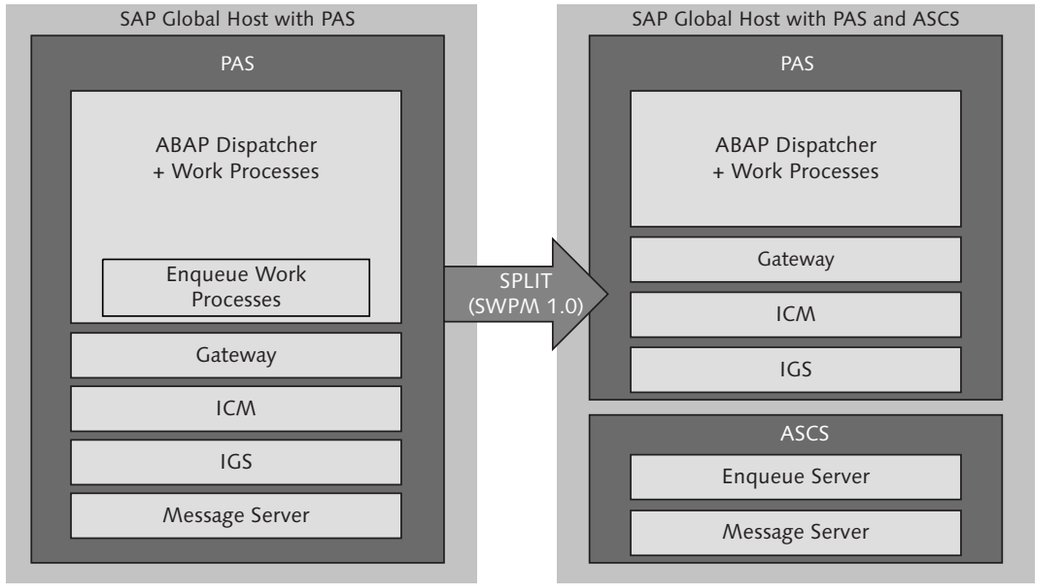


**Figure 3.1** ABAP SAP Central Services Instance

Follow these instruction to split the enqueue server from the central instance:

1. Using the SWPM, launch the application on the PAS.

2. In the directory tree, choose <Your SAP NetWeaver Release> • <Your Database> • Additional SAP System Instances • Split Off ASCS Instance from Existing Primary Application Server Instance (see Figure 3.2).

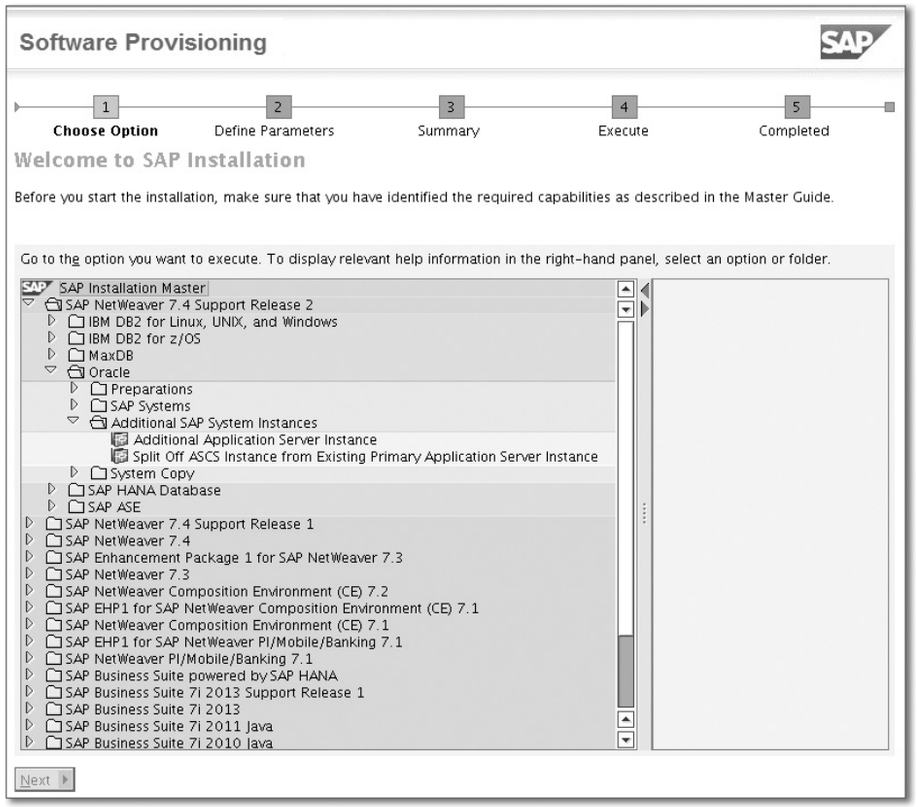


**Figure 3.2** Guided Procedure to Split Off the ABAP SAP Central Services Instance

3. Choose Typical as the Parameter Mode when prompted.
4. Enter the path to the SAP system profiles in the Profile Directory field shown in Figure 3.3, and click Next.
5. Input the Database ID (DBSID) and the Database Host, and click Next.





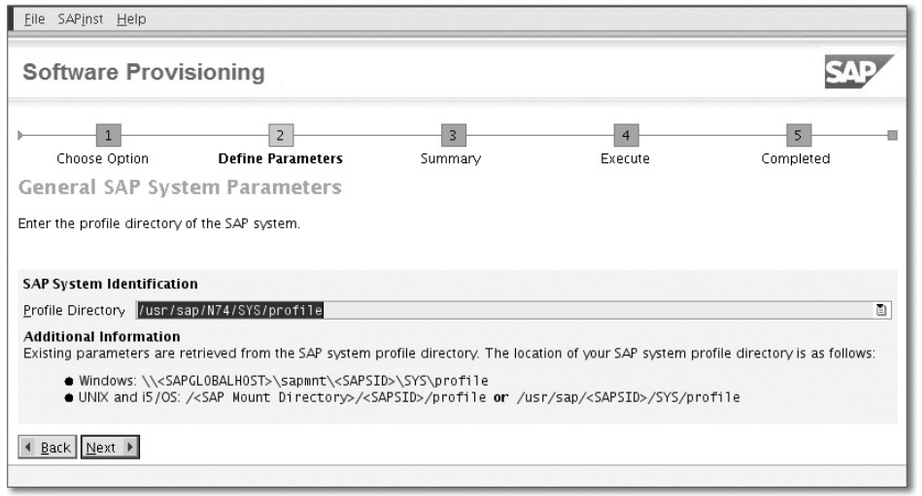


**Figure 3.3** Reading the Profiles to Validate the Process

6. Define the path to the kernel media (which can be downloaded from the SAP Support Portal at *http://support.sap.com/swdc*), as shown in Figure 3.4.

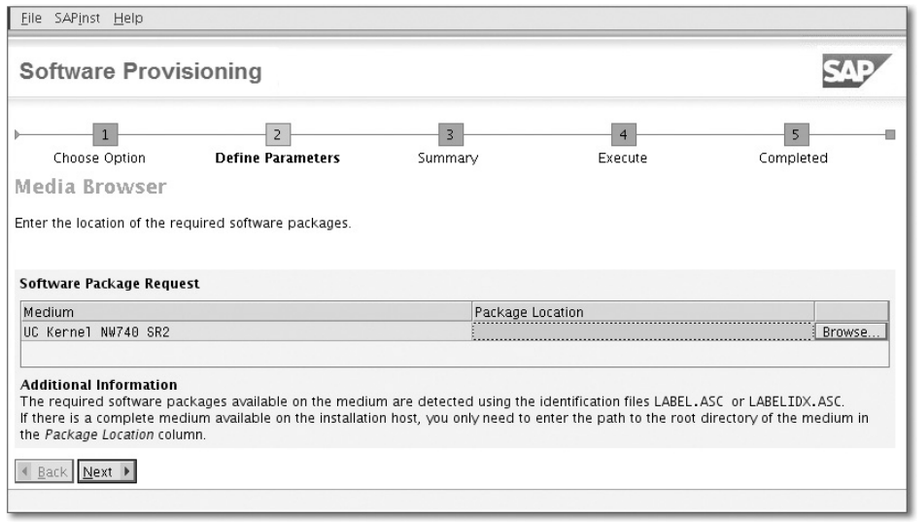


**Figure 3.4** Entering the File Path for the Kernel Media

7. The ports for the ABAP message server are read automatically. Validate, and click NEXT.
8. On the SUMMARY step, verify all input. Clicking SHOW DETAIL will give you all parameters used. Then click NEXT, as shown in Figure 3.5.

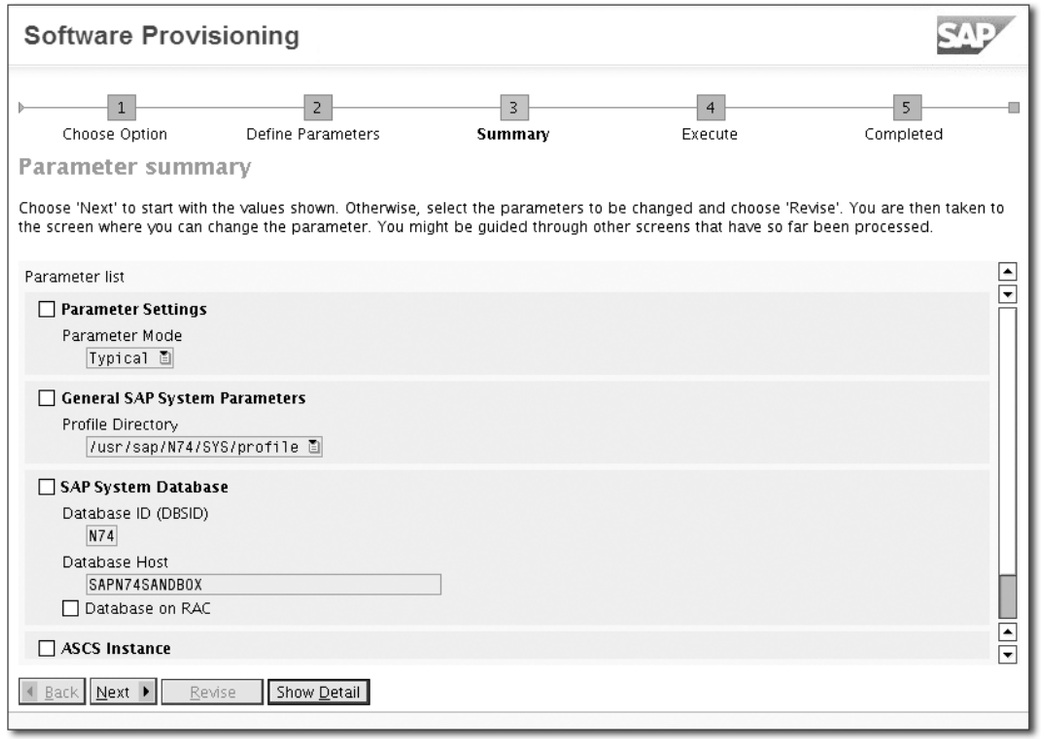


**Figure 3.5** Summary of All Parameters

9. The installer processes the information. During this time, the application server instances are restarted with new parameters.
10. When prompted, confirm the completion message.

[»] **Note**

For information on the process of splitting the ASCS instance and additional settings, see SAP Note 2119669 – How to Split the ASCS from Primary Application Server (PAS).





## 3.2 Dual Stack Split

SAP no longer supports the update of dual stack systems with the SAP Business Suite. Dual stack systems are defined as a single SAP instance with both an ABAP AS and a Java AS sharing a single system ID and database. This increases maintenance of the system and decreases system performance. Fortunately, SAP provides the SWPM as part of the software logistics toolset (SL Toolset). This program provides a menu-driven wizard tool to export the Java stack, install a new Java stack, and then uninstall the Java stack from the ABAP system. Several options are available to you in terms of where the new Java instance resides. You can install the new Java instance on the same host(s) as the ABAP instance, or use this opportunity to dedicate a host for the Java instance, as shown in Figure 3.6. This answer depends on your current Java instance usage and current hardware sizing. Installing the Java stack on the same host requires additional resources because the new instance will have its own set of binaries, database, and lifecycle.

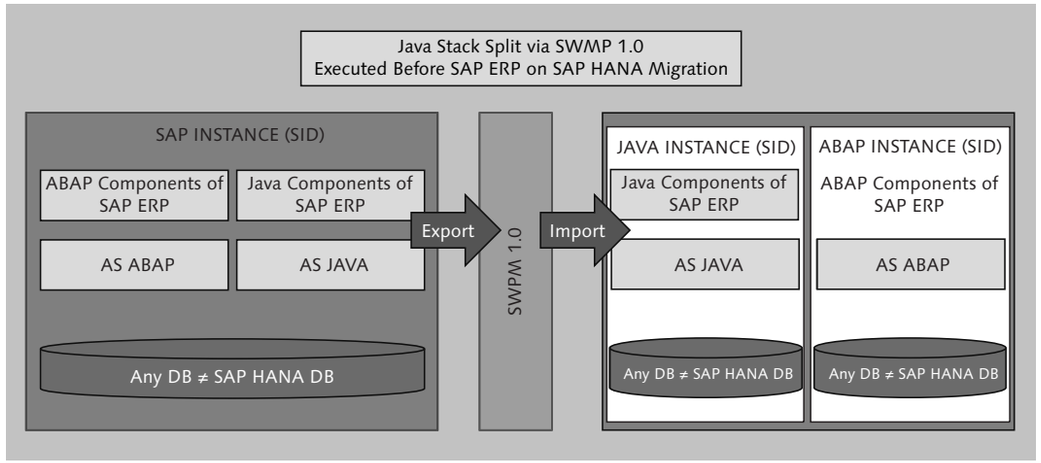


**Figure 3.6** Dual Stack Split Process Overview

It's worth noting that the Java stack split creates a completely separate Java instance with its own User Management Engine (UME). This means your SAP ERP (ABAP) users who previously could access the Java stack as part of SAP ERP will no longer have a user ID and password to the Java instance. However, the Java instance UME can be reconfigured to use the ABAP system as the UME, meaning

that the Java stack user management database can continue to reside within the ABAP stack.

[»] **Note**

For the latest information on the SL Toolset, which includes SWPM and the Software Update Manager (SUM), check out the SAP Service Marketplace at *http://service.sap.com/sltoolset*. This site is dedicated to release information, installation guides, how to start/stop SAPINST, and direct access to downloads.

### Executing Dual Stack Split

To execute the dual stack split, follow these steps:

1. After downloading the latest SWPM from SAP Service Marketplace, log on to the PAS with an administrative account.
2. Start SAPINST, and navigate via the menu to DUAL-STACK SPLIT. The first task you'll execute is EXPORT JAVA STACK, as shown in Figure 3.7.

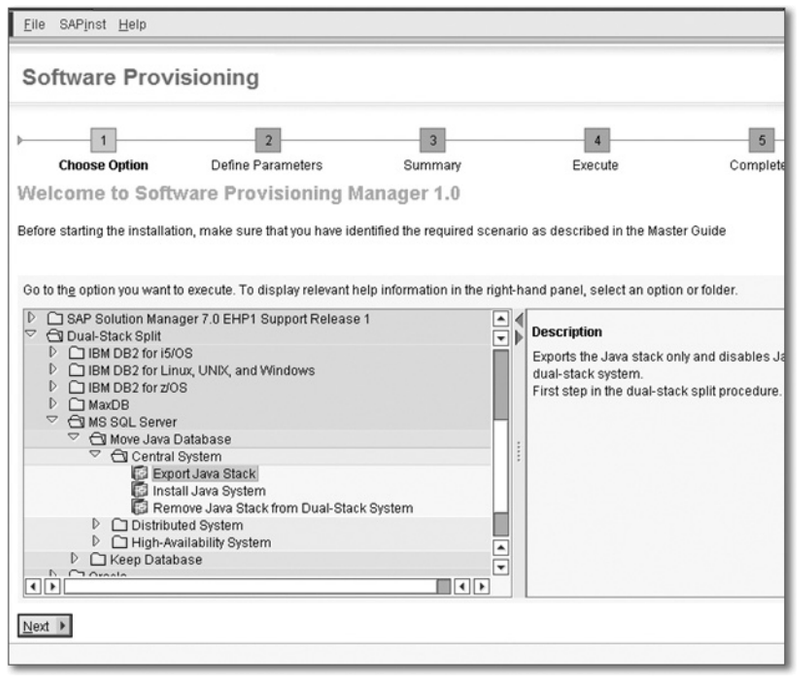


**Figure 3.7** Initial Menu in Software Provisioning Manager





3. Select the menu item, and click NEXT.
4. When prompted, enter the path to the SAP PROFILE DIRECTORY, and click NEXT.
5. Enter the PASSWORD OF SAP SYSTEM ADMINISTRATOR when prompted, and click NEXT.
6. Select the correct DATABASE ID (DBSID), and click NEXT.
7. On the next screen, select the EXPORT LOCATION to store the Java export, and leave MANUALLY STOP SYSTEM checked, as shown in Figure 3.8. Typically, 10–20GB will suffice.


**Figure 3.8** Choosing the Export Location

8. Select the option to let the SUM disable the Java stack using the profile parameter `rdisp/j2ee_start=0`. This doesn't stop the system at this point, but it adds the required parameters to disable the Java stack from starting when you restart services. Check the DISABLE APPLICATION SERVER JAVA box, and then click NEXT.
9. On the PARAMETER SUMMARY screen (see Figure 3.9), review all parameters, and click NEXT to begin the export process. You'll need to manually stop SAP services when your team gives you permission.

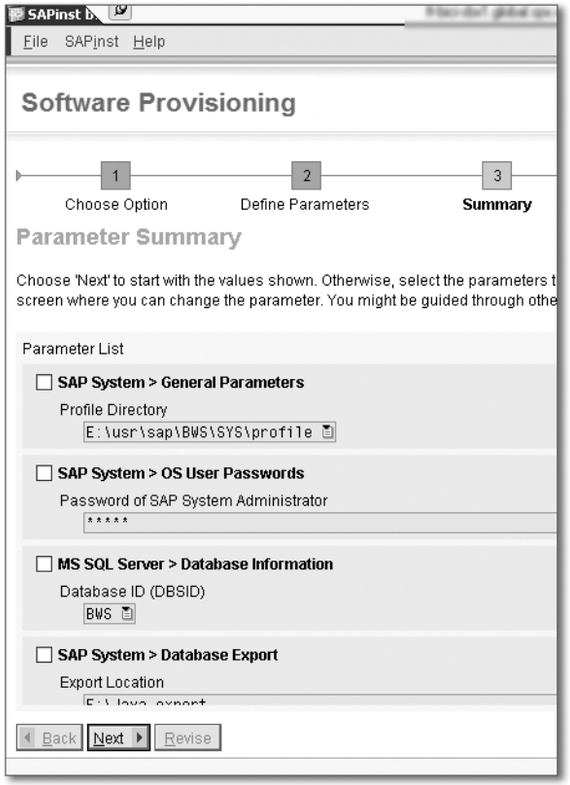


**Figure 3.9** Parameter Summary Screen in Software Update Manager

10. At this point, the Java export occurs and should not be restarted until you install the new Java instance.
11. When prompted, stop SAP services, and click OK (see Figure 3.10). Note you only need to stop Java services within the dual stack system using Transaction SMICM.
12. Click OK. The export process continues, and the application server downtime begins for the dual stack split.
13. Verify the results on screen when presented with the completion message (see Figure 3.11).
14. Restart the SAP services using the command `Stop SAP`, or use the Microsoft Management Console (MMC).





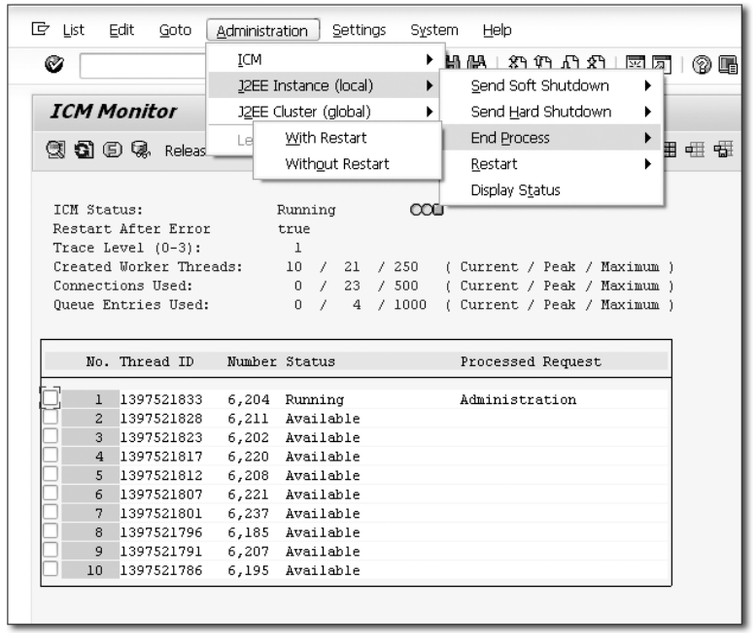


**Figure 3.10** Stopping the Java Instance

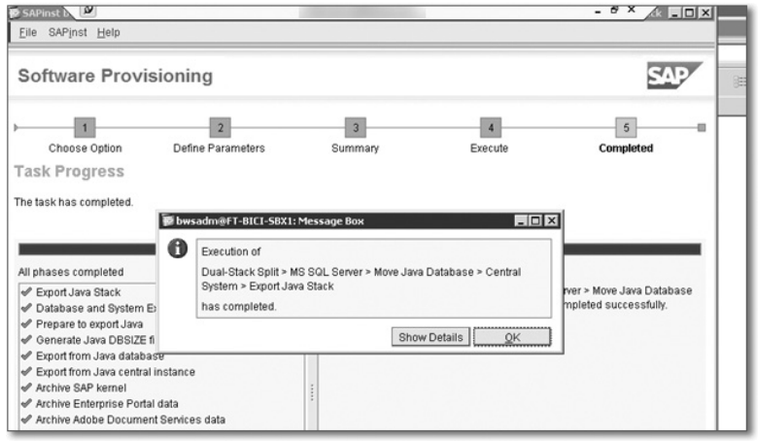


**Figure 3.11** Confirmation Message after the Export Completes

15. Return to the SWPM by running SAPINST again. Navigate within the SWPM menu to DUAL-STACK SPLIT • <DATABASE> • MOVE JAVA DATABASE • CENTRAL SYSTEM • INSTALL JAVA SYSTEM (see Figure 3.12).

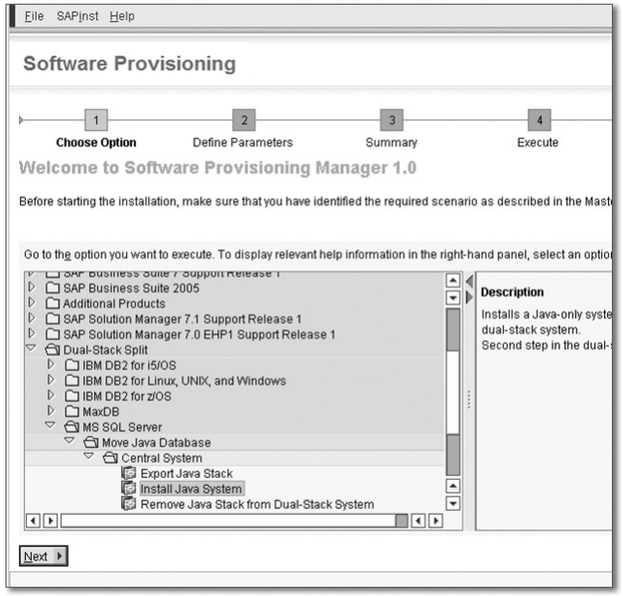


**Figure 3.12** Menu Path to Install New Java System

16. When prompted, enter the PACKAGE LOCATION of the Java export created in an earlier step (see Figure 3.13). Click NEXT.

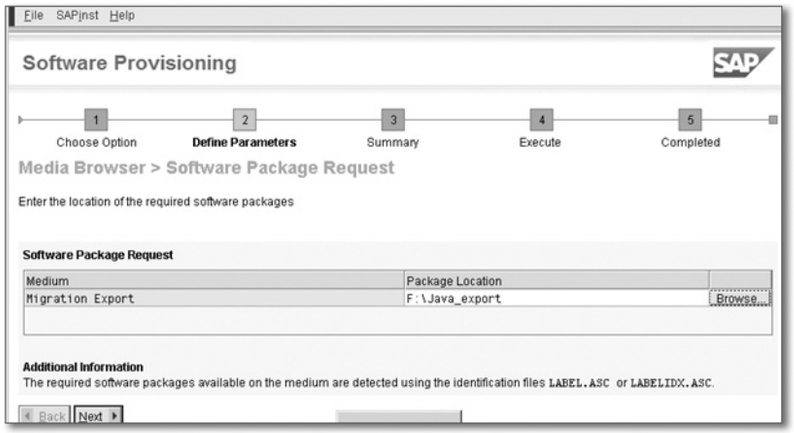


**Figure 3.13** Enter the Java Export Directory

17. Select TYPICAL for the PARAMETER MODE, and click NEXT.





18. Enter the new Java System ID (SID) and the location for the new installation directory, which will contain the executables for the Java application server, as shown in Figure 3.14.

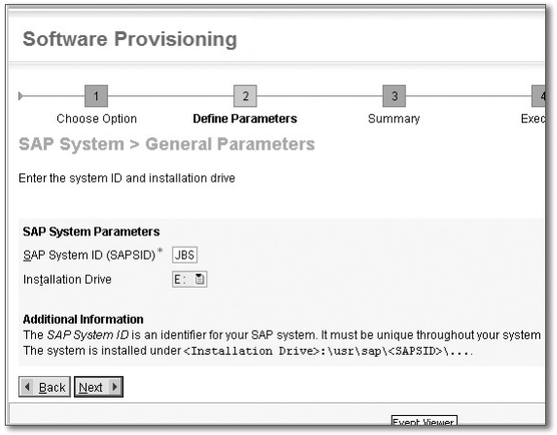


**Figure 3.14** SAP Paramters Entry Screen

19. Assign a MASTER PASSWORD, and click NEXT. (Note that the default password for system accounts will be the master password.)

20. Input the SAP SYSTEM ADMINISTRATOR and SAP SYSTEM SERVICE USER passwords, and click NEXT (see Figure 3.15).


**Figure 3.15** Setting Passwords

21. When prompted to select the Database to Use, click Create New Database (DBSID = SAPSID), and click Next.
22. For the Expected System Size section, select the system size based on the number of processor cores assigned to this instance (see Figure 3.16). Click Next.

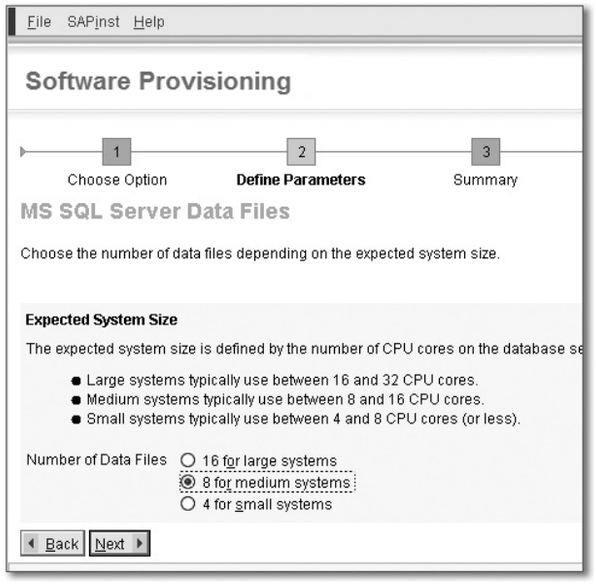


**Figure 3.16** Selecting the System Size for Database Data Files

23. In the Software Package Request section, enter the location of the Java Component media in the Package Location field (see Figure 3.17), which can be downloaded from *https://support.sap.com/swdc*.
24. When prompted, enter the administrator password for the previously exported Java system, and click Next.
25. Enter the location of the SAP Cryptographic Library to be used during the installation (see Figure 3.18). This library (or SAR file) can be downloaded per instructions from SAP Note 455033.





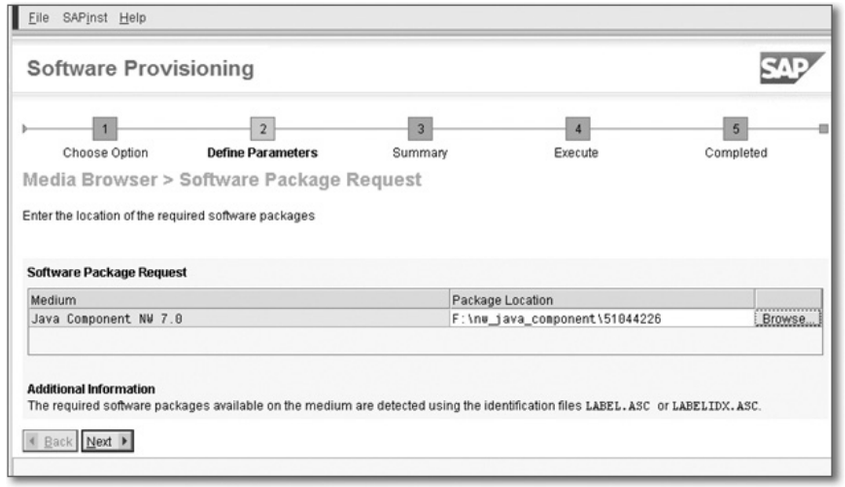


**Figure 3.17** Entering the Directory Containing Downloaded Media

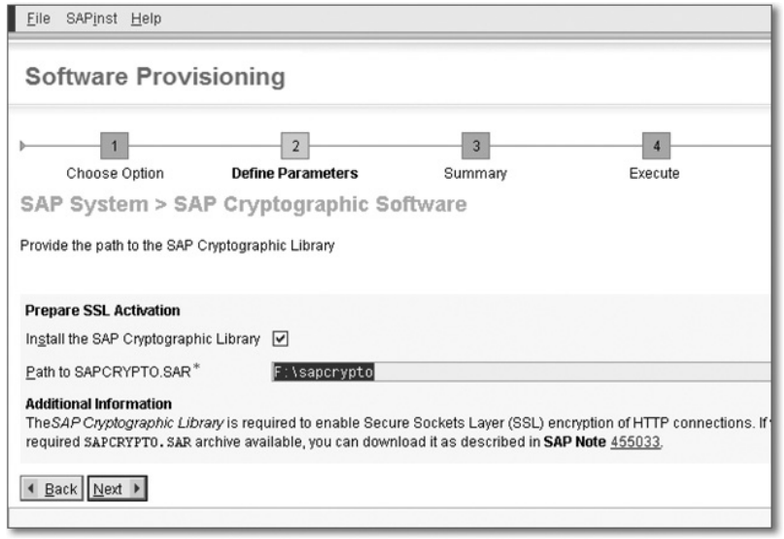


**Figure 3.18** Installing the SAP Cryptographic Library

26. On the PARAMETER SUMMARY screen, review the selected criteria and choose NEXT (see Figure 3.19).

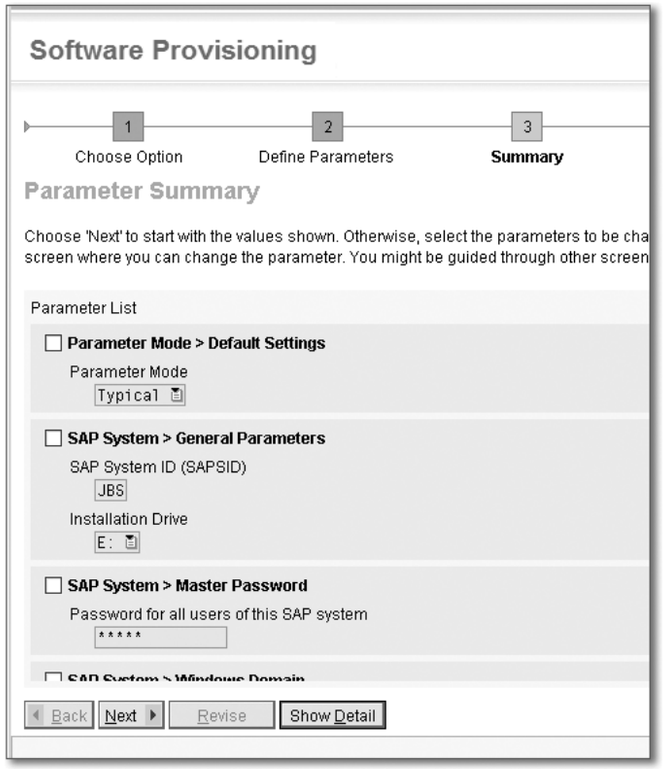


**Figure 3.19** Software Povisioning Manager Summary Screen

27. When prompted, verify the ABAP stack is running, and click OK (see Figure 3.20).

28. Verify the results when prompted with a completion message (see Figure 3.21).





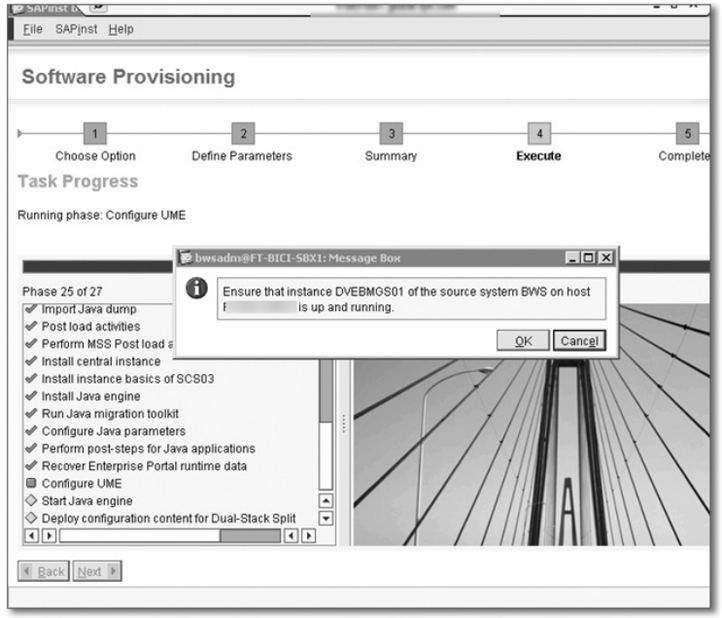


**Figure 3.20** Confirming the ABAP System Is Running

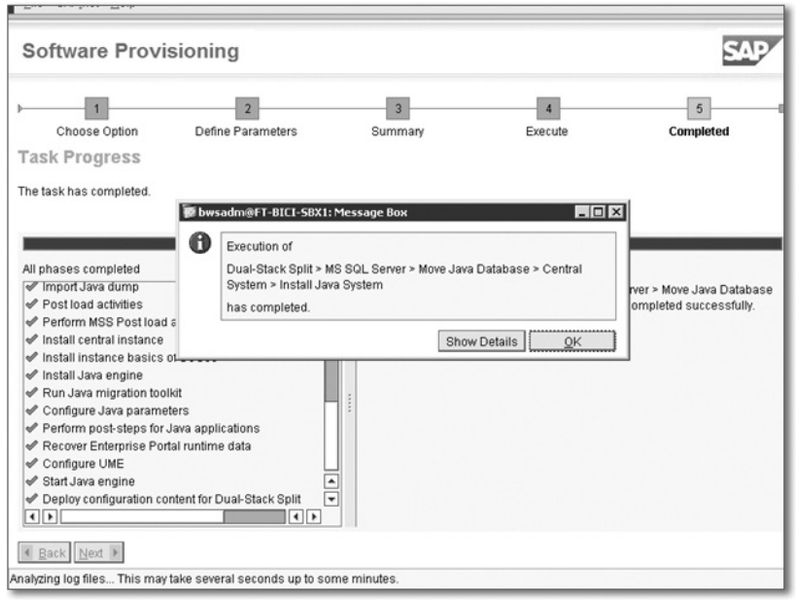


**Figure 3.21** Completion Message from Software Provisioning Manager

29. Return to the SWPM, and run SAPINST. Navigate within the SWPM menu to Dual-Stack Split • <Database> • Move Java Database • Central System • Remove Java Stack from Dual-Stack System (see Figure 3.22).

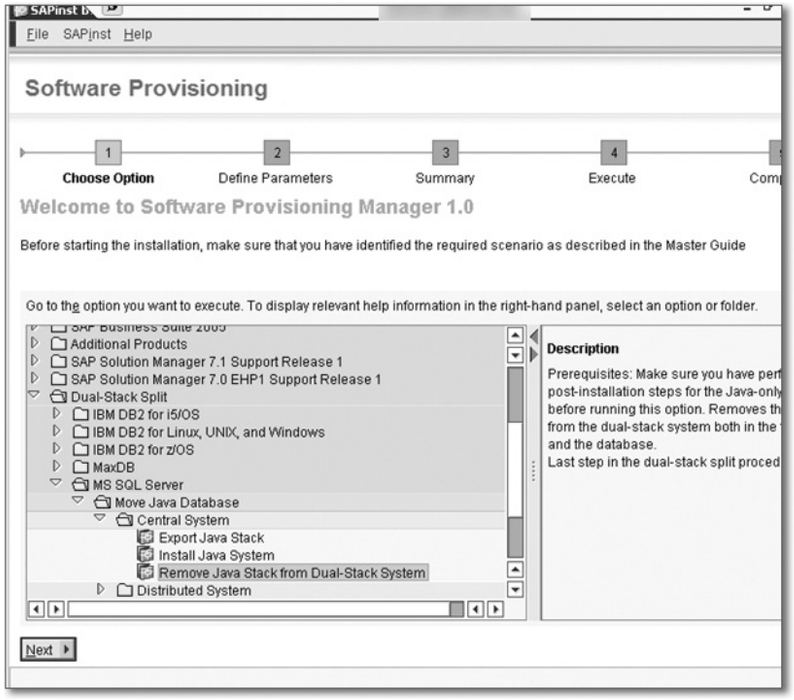


**Figure 3.22** Final Step in the Dual Stack Split

30. Enter the path to the ABAP system SAP Profile Directory (see Figure 3.23). Click Next.
31. Enter the Password of SAP System Administrator for the ABAP stack.
32. Next select the Database ID (DBSID) for the ABAP system. Click Next.
33. Verify the selection criteria on the Parameter Summary screen, and click Next to begin the removal process (see Figure 3.24).





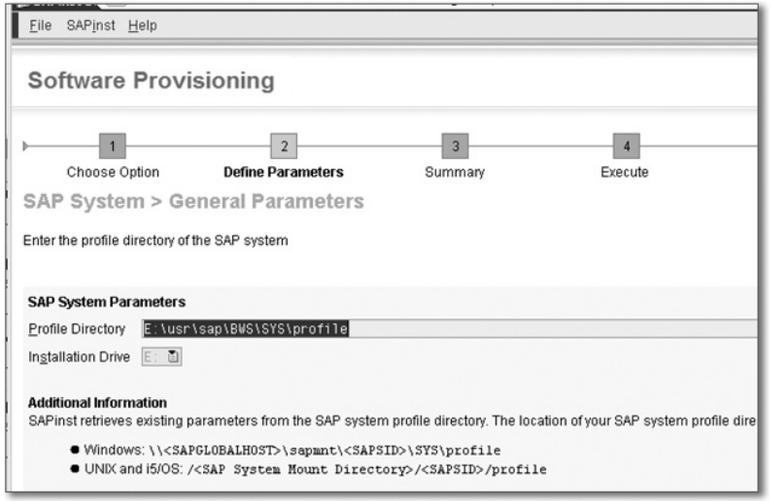


**Figure 3.23** Entering the Path to the ABAP System Profiles

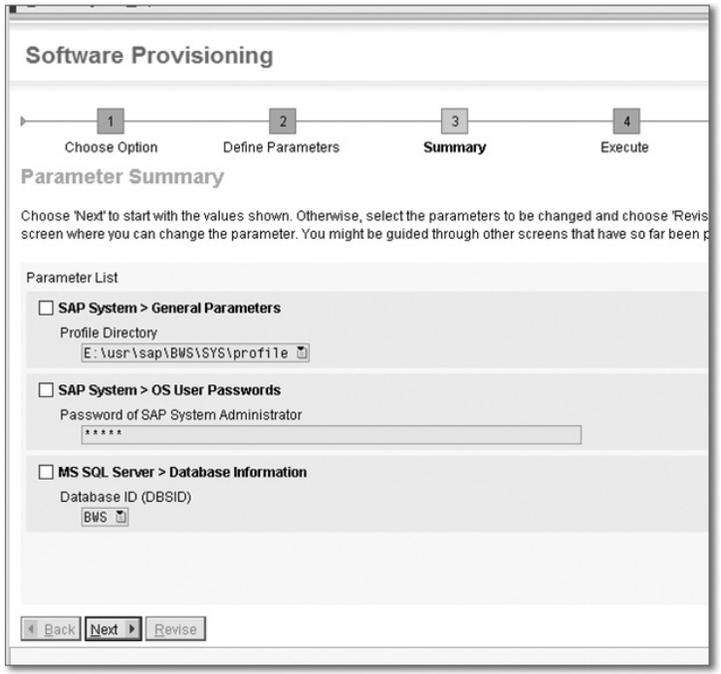


**Figure 3.24** Parameter Summary Screen

34. On completion of the task, verify the results when prompted (see Figure 3.25). Click OK.

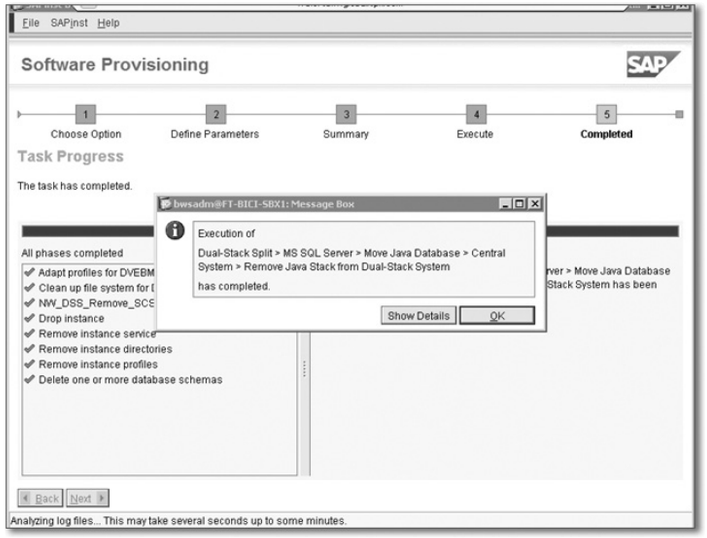


**Figure 3.25** Verification of Completion Message from Software Provisioning Manager

35. Confirm the new Java system is running by executing `StartSAP` as the new Java <SID>ADM user account or by opening MMC to confirm the new instance is running (see Figure 3.26).

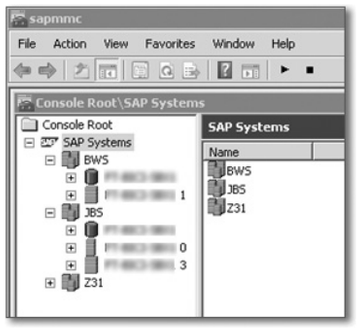


**Figure 3.26** New Java System Displayed and Running

Separating the ABAP and Java stack is a requirement that must be completed prior to starting the migration to SAP Business Suite on SAP HANA. These procedures





can be executed weeks or months in advance of the migration project. Keep in mind that the screenshots and process in this section assumed no changes in the hardware. The Java instance was separated but remains on the same application servers. This approach limits the amount of change so that testing can be focused at the application layer. If you so choose, you could split the Java stack and move the application server to new hardware or a virtualized OS.

## 3.3 Sizing

Is SAP sizing more art than science? Some consultants think so. SAP has provided clear guidelines on most if not all components, and while there are spot solutions with fuzzy sizing guidelines, the primary components are well documented. SAP also assumes additive sizing. Do you want to run the ABAP application server on the same host (appliance) as the SAP HANA database? Yes, it's supported. And yes, you add the sizing requirements together (ABAP + SAP HANA) to get the total size required. Keep in mind that sizing isn't done in a silo, and it's not done just once. During your SAP Business Suite on SAP HANA migration project, you're going to learn a lot, and you'll have performance data from all the regression tests. Ideally, you'll mitigate the risk of a bad sizing by running a load test before you go-live.

In the beginning, you'll use the ABAP report provided by SAP to perform a SAP Business Suite on SAP HANA sizing from your production system. This provides an estimate for main memory requirements of the SAP HANA database but won't provide an estimate of SAPS (processing power) required for the application servers. At the start of your project, your sandbox will likely have one PAS connected to a SAP HANA appliance. This might be true for development as well because the load will be small. As your project moves forward, we typically see customers with QA environments that closely match their production environments. This QA landscape that closely resembles production is the ideal environment for a load test. Let's explore some of the sizing methods available to you as a customer. And remember—most hardware partners have an SAP sizing competency center that can help you as well.

**What are SAPS?** [«]

SAP Application Performance Standard (SAPS) is a hardware-independent unit of measurement that defines the performance of a system operation in the SAP environment. It was created to independently measure performance across different types or processors.

### 3.3.1 Ballpark Sizing

There is a formula you can use very easily to calculate the approximate size of the SAP HANA appliance you'll need for your SAP Business Suite on SAP HANA migration. It first requires the size of your existing database, which you can retrieve using Transaction DB02. The recommendation from SAP is that you take half of the size of your disk-based database, include a safety buffer of 20%, and add 50GB for the repository, stack, and other services. Let's look at an example:

1. Log on to your SAP ERP system with a valid user ID and password.
2. Run Transaction DB02, and note the database size displayed (see Figure 3.27).

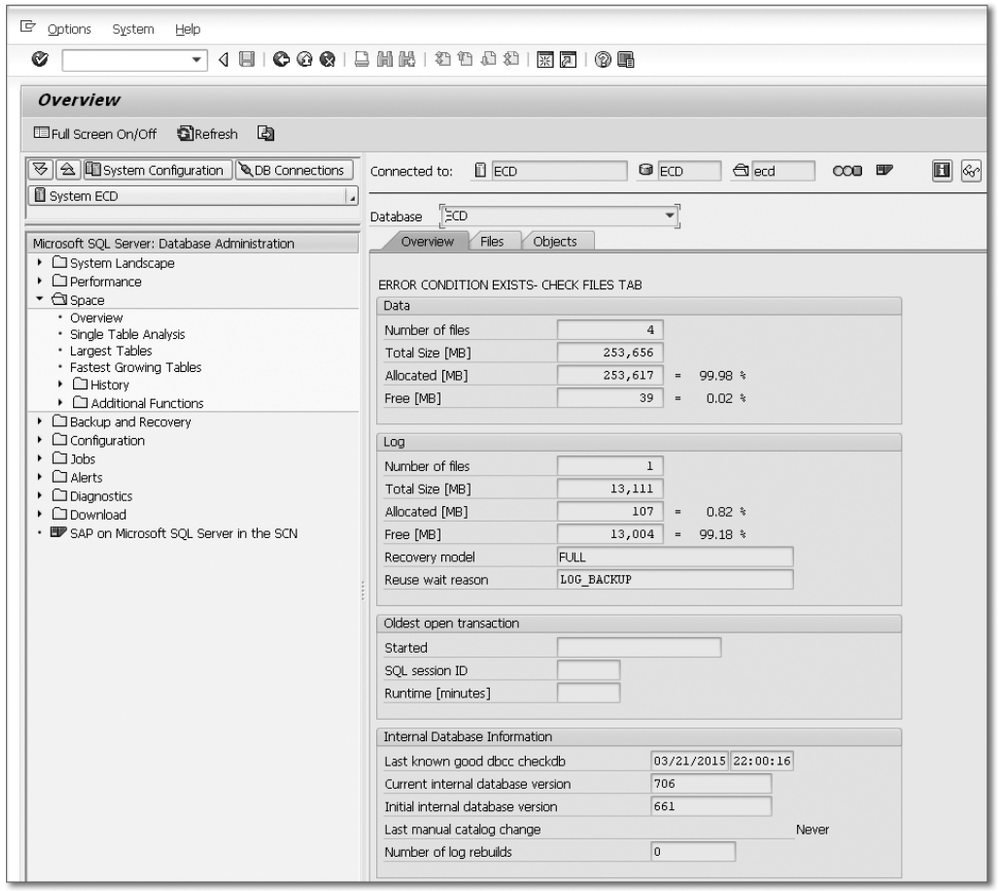


**Figure 3.27** Database Size of 253GB





3. Using SAP's formula, you can calculate the approximate size of your SAP HANA database. In this example, this roughly calculates a SAP HANA database of 201GB:

   *HANA DB Size = (253GB ÷ 2 × 1.2 + 50GB)*

### 3.3.2 SAP Sizing Report for SAP Business Suite on SAP HANA

If your system is already running SAP ERP 6 EHP 7 with the latest versions of the ST-PI ABAP add-on component, then the SAP Business Suite on SAP HANA sizing report already exists in your system. SAP has provided updates to the program, which can be downloaded to your system by implementing SAP Note 2175150 – Suite on SAP HANA Memory Sizing Report. After you've downloaded and implemented the SAP Note, follow these instructions:

1. Log on to your SAP ERP system with a valid user ID and password.
2. Run Transaction SE38, and input the program name "/SDF/HDB_SIZING". Click Execute.
3. By default, you'll leave the List of Tables fields blank because you want to get the size for the entire system.
4. Uncheck Sizing for SAP Simple Finance 2.0 (alternate name: Sizing for SAP S/4HANA Finance) because you're sizing for the migration only. Your team will need to size for SAP S/4HANA Finance later.

**Note** [«]

If you need to size SAP HANA Enterprise for specific tables in SAP ERP, you can use the same program. If your organization wants to implement the SAP HANA sidecar scenario, you can specify the exact tables to be included and get a good sizing for those specific objects.

5. Under Technical options, enter the Number of parallel dialog processes you want to use. Leaving this with the default of 1 will result in extremely long runtimes.

6. After you make your selections, select PROGRAM • EXECUTE IN BACKGROUND, as shown in Figure 3.28.

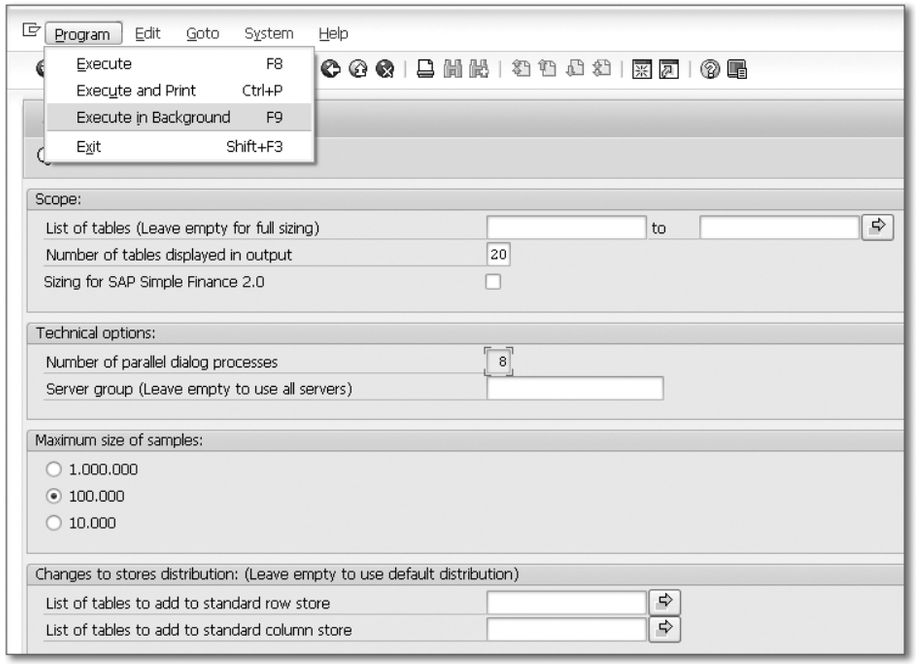


**Figure 3.28** Example Options Entered into the SAP Business Suite on SAP HANA Sizing Report

7. After the background job completes, you'll see the created spool file.
8. In Figure 3.29, you can see the as-is requirement for this example system on SAP HANA will be 135GB of main memory. If we add in a 20% safety buffer, then we arrive at a requirement of 162GB of memory. That is about 40GB less than our ballpark estimate for the same system.





***Graphical display of spool request 29640 in system ECD***

Settings... Graphical Graphic Without

RESULTS OF SUITE ON HANA SIZING IN GB

Based on the selected table(s),the anticipated memory requirement is:

| | |
|---|---|
| Column store data | 25.1 |
| + Row store data | 14.9 |
| + Work space | 40.0 |
| + Hybrid LOB stored on disk | 28.0 |
| + Fixed size for code, stack and other services | 50.0 |
| = Anticipated total memory requirement considering 20% records stored on disk are cached in memory | 135.6 |

| | |
|---|---|
| Sizing report: | /SDF/HDB_SIZING |
| SID | ECD |
| NW release: | 740 SP 8 |
| Date of analysis: | 06/23/2015 |
| Version of the report: | 44 |
| Selected accuracy: | M |
| Number of work processes used: | 06 |
| Type of analyzed database: | MSSQL |
| Number of tables successfully analyzed: | 95,946 |
| Number of tables partially analyzed: | 0 |
| Number of tables with error: | 0 |

| LARGEST COLUMN STORE TABLES | ESTIMATED HANA MEMORY SIZE IN GB | ESTIMATED RECORD COUNT |
|---|---|---|
| EDID4 | 3.5 | 73,210,554 |
| EDIDS | 2.9 | 55,048,214 |
| WBCROSSGT | 2.5 | 49,248,714 |
| BSEG | 1.4 | 15,017,812 |
| DOKTL | 1.3 | 34,145,251 |
| FAGLFLEXA | 1.0 | 15,017,813 |
| BSIS | 1.0 | 15,017,688 |
| COEP | 0.9 | 15,017,713 |
| REPOSRC | 0.9 | 4,125,869 |
| EDIDC | 0.6 | 18,458,581 |
| BALDAT | 0.5 | 1,605,662 |
| SCPRVALS | 0.4 | 10,107,814 |
| SCPRSVALS | 0.4 | 9,044,524 |
| E071K | 0.4 | 8,496,894 |
| BKPF | 0.4 | 7,508,939 |
| /SLOAP/CLU_STORT | 0.4 | 232,248 |
| COBK | 0.3 | 7,508,994 |
| E071 | 0.2 | 8,354,184 |
| SEOSUBCOTX | 0.2 | 4,126,105 |
| OCSCMPLOBJ | 0.2 | 2,572,931 |

**Figure 3.29** Example Output from the SAP Business Suite on SAP HANA Sizing Report

### 3.3.3 SAP Quick Sizer

The SAP Quick Sizer is another tool you can use to accurately size your SAP Business Suite on SAP HANA environment. There are two means of performing a sizing within the Quick Sizer: user-based sizing and throughput-based sizing. The latter requires more time and effort. User-based sizing is a count of users by activity type (more on that shortly). Throughput-based sizing focuses on the quantity of documents within your systems and the frequency at which they occur. Getting this information will take time, but it can provide an extremely accurate sizing estimate. However, as we've said before, the only way to mitigate the risk of a potentially bad sizing is to perform a performance test of the target landscape before go-live. If your team doesn't have the capacity to perform a throughput sizing, you can mitigate this risk with a performance test. Our goal with this section is to provide you with enough detail to understand the effort required and where to find more information on the Quick Sizer.

A user-based sizing will give you the quickest time to a number, which you can then provide to a hardware vendor. User-based sizing breaks down users into three categories by user activity type because not all users are created equal. The Quick Sizer also lets you enter user counts by time period. If your organization operates in multiple time zones, this feature will let you input user activity at varying times in the day. User-based sizing also has a field to capture data retention period by module. In the following examples, we entered a value of 84 months or 7 years. This is only an example, and ideally your organization has some data retention policy you can reference when completing these steps. Here's a summary of the user activity types you can use in user-based sizing:

- **Low-activity user**
  This user is an informational or executive user in SAP ERP. These users perform 10 dialog steps within an hour.
- **Medium-activity user**
  This user is someone whose daily job duties primarily occur in SAP ERP, for example, accountants, clerks, shop floor workers, and so on. These users perform 120 dialog steps an hour or one click in SAP ERP every 30 seconds during work hours.
- **High-activity user**
  This is a heavy user of the system, for example, customer service agents, data entry users, or system administrators. These users perform 360 dialog steps in an hour or one click every 10 seconds in SAP ERP.





You can easily find the activity level of users within your system by following these steps:

1. Log on to SAP ERP with a valid user ID and password.
2. Run Transaction ST03N.
3. From the upper-left area, expand the node Workload • Total • Month, and double-click the last full month of data.
4. Double-click Workload Overview under Analysis Views in the lower-left area.
5. Select the User tab to see the user types by activity, as shown in Figure 3.30.

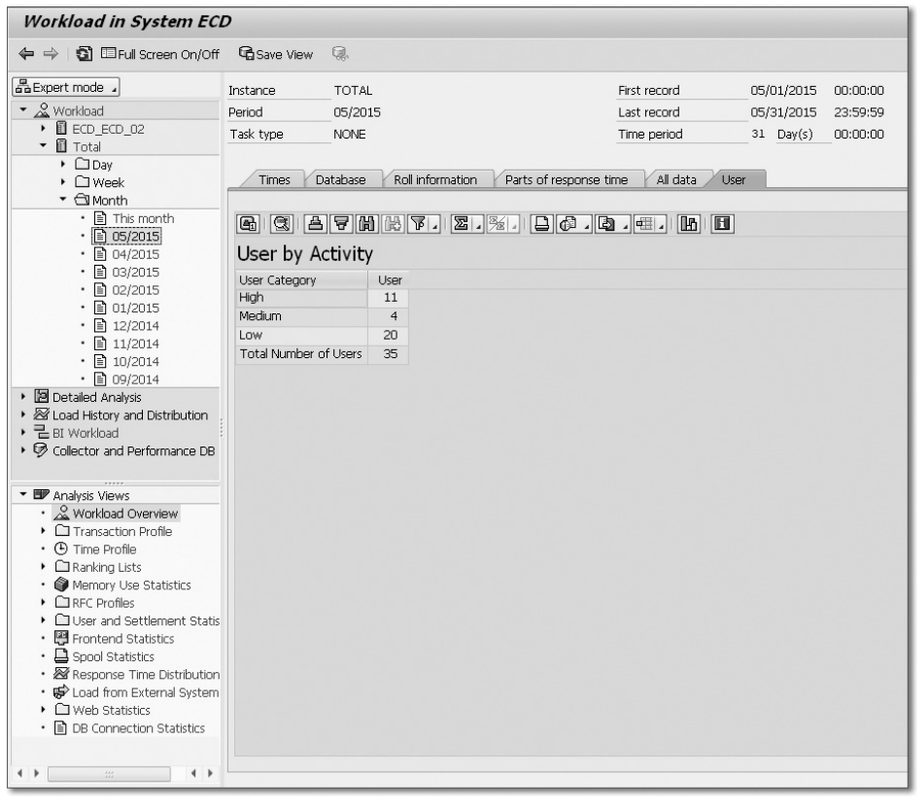


**Figure 3.30** Users Listed by Activity Type

Now, you'll notice the rules give you the activity type by user but don't break it down by functional area. Your team will have to determine this. You can look at

logons for the month analyzed to develop a list of users that you've assigned to a primary functional area. You know that users typically don't work within one module, so you need to identify the primary functional area they are responsible for.

When accessing the Quick Sizer, you'll notice your customer number and a project name displayed. These are important identifiers because after your sizing estimation is complete, you can share your Quick Sizer report directly with certified SAP hardware providers or other individuals in your company. The project name can be any text value that you associate with the project activity (i.e., SAP HANA Migration).

You'll also notice in the Quick Sizer that your project will have a status value. When setting your project to GoingLive, this will enable integration between your SAP Enterprise Support service called the GoingLive Analysis with the Quick Sizer project. SAP Support will essentially review your sizing project and then compare it to the actual hardware you've deployed. This ensures that your actual system will meet the needs of the sizing you estimated.

### Accessing Quick Sizer

To access Quick Sizer, follow these steps:

1. Log on to the SAP Support Portal via the direct link to the SAP HANA version of the Quick Sizer at *http://service.sap.com/hanaqs.*
2. On the initial screen, enter a PROJECT NAME, and click CREATE PROJECT, as shown in Figure 3.31.

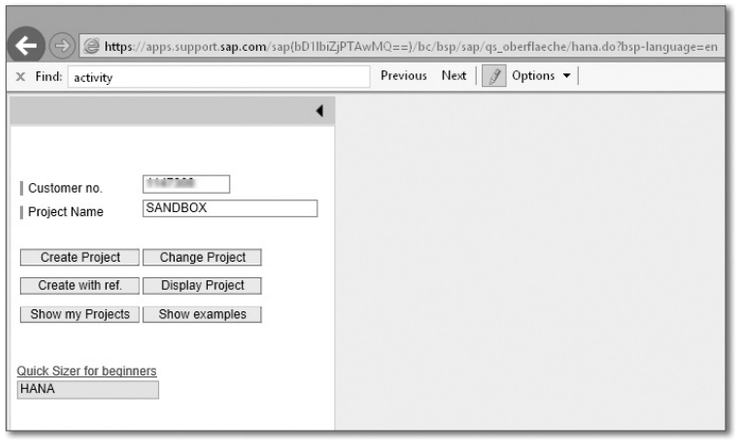


**Figure 3.31** Example Entry Screen for the SAP HANA Version of the Quick Sizer





3. Initially, the PROJECT SANDBOX screen appears with several helpful items displayed (see Figure 3.32). Click HARDWARE VENDORS to get contact information for the SAP sizing competency center for each vendor.
4. Click QUICK SIZER TOOL DOCUMENTATION to access the guide to using the tool.
5. In the NEW SYSTEM/SYSTEM EXTENSION section, you must decide between NEW SAP BUSINESS SOLUTION/SOFTWARE COMPONENT and SAP BUSINESS SOLUTION/SOFTWARE COMPONENT EXTENSION. If you size as a new solution, your hardware can tell you what existing hardware will support what you've sized.
6. The PLATFORM AND COMMUNICATION, SYSTEM AVAILABILITY, and NETWORK INFRASTRUCTURE sections are optional. They help provide the hardware vendor and SAP with more background information.


**Figure 3.32** Project Information Screen in the Quick Sizer

7. If you look to the left of your screen, you can see a navigation tree with multiple components under SAP BUSINESS SUITE ON SAP HANA.
8. The items you select will depend on the scope within your organization's SAP ERP environment. In the example in Figure 3.33, there is data for a user-based

sizing with SAP S/4HANA Finance, Logistics Execution, Product Dev & Execution, and Sales & Service.

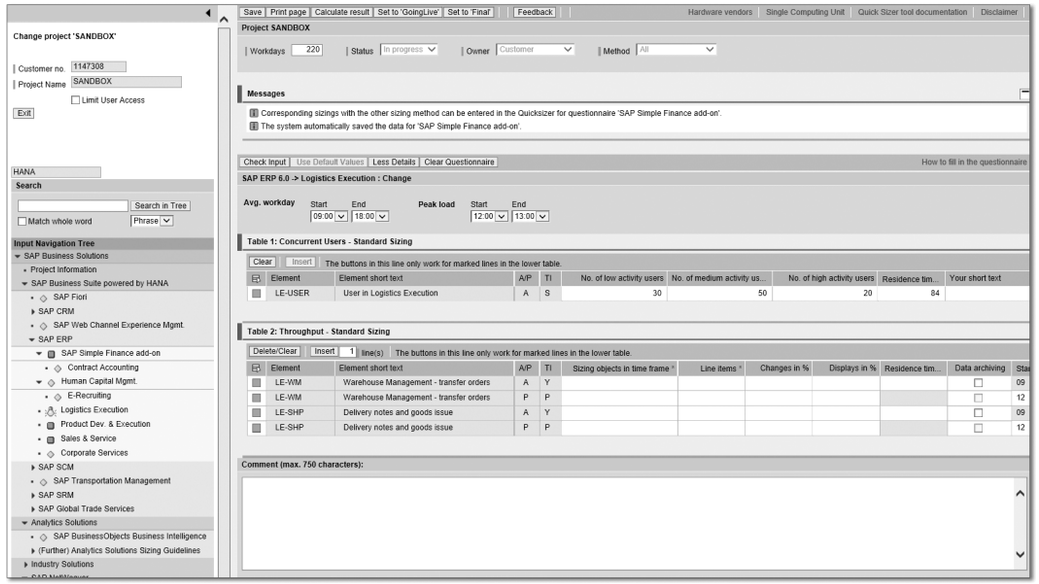


**Figure 3.33** Example User-Based Sizing for SAP Business Suite on SAP HANA

9. In Figure 3.33, you can see that we entered a total of 100 Logistics users, with a data retention of 84 months for this example. We entered similar data for the other SAP ERP components to simulate a sizing for a system with 600 active users.
10. Click on Calculate result at the top of the screen to see a summary of the sizing (see Figure 3.34).

In Figure 3.34, we used the default of 220 workdays per year, with the average workday running from 9AM to 6PM, and peak load occurring from 9AM to 11AM and 1PM to 2PM. Based on the 600 users we entered, with a retention period of 84 months for SAP ERP, the Quick Sizer believes we'll need a SAP HANA appliance with essentially 1TB of capacity to support both our user load and data volumes. We also get a SAPS number for the application server in the sizing summary. The Quick Sizer doesn't know if our ABAP AS will be on the same hardware as SAP HANA or separate. We would add these values together if they will be installed on the same host.





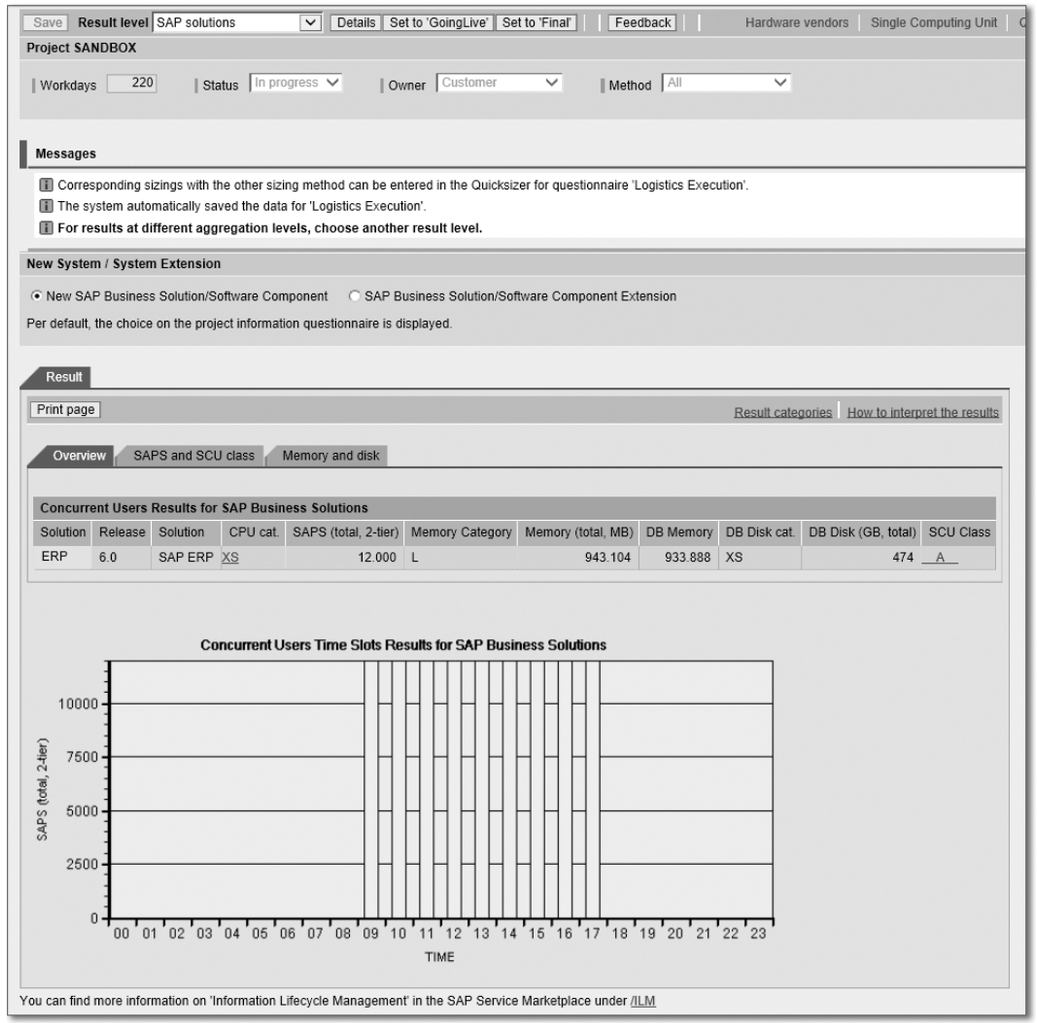


**Figure 3.34** Example Sizing Results Displayed

Sizing your new SAP Business Suite on SAP HANA system using the Quick Sizer can get very complicated if you move into the throughput-based sizing method. Using this sizing method, you need to capture or estimate specific activity within your SAP ERP system. If you believe you require this level of sizing, we recommend engaging a certified SAP sizing professional either through SAP or one of its certified partners.

In this section, we described how to perform initial sizing estimations, how to execute the ABAP report for SAP Business Suite on SAP HANA sizing, and how to perform the user-based sizing method using the Quick Sizer. SAP and its certified hardware partners are available to support your organization through this process. Uncertainty as to what hardware you need to migrate SAP Business Suite on SAP HANA can be easily remedied by using the solutions described in this chapter and engaging your hardware partner early in the process. SAP sizing tools exist to ensure success at your go-live. This isn't a process you'll only execute once during the project. As you continue to learn during the upgrade and migration process, you'll need to add reminders throughout your project to revalidate the sizing that was completed. And, don't forget, you can mitigate risk due to a bad sizing by running a performance test during your project.

[»]

**Note**

For even more information on SAP sizing, training events, and how-to guides, see the SAP Service Marketplace at *https://service.sap.com/sizing*. Click on the SAP HANA Quick Sizer link to access the SAP Business Suite on SAP HANA sizing tool.

## 3.4 Backup and Recovery

When SAP HANA was initially released and touted as an in-memory database, the IT person in all of us wondered how data in-memory is handled when the system crashes. Rest assured, SAP thought of this well before releasing SAP HANA to the public. SAP HANA is an in-memory database that also uses persistent storage. This allows you to backup and recover the database just like any other database on the market. You have graphical, command-line, and third-party backup tools that enable you to recover when the power is cut, a disk drive fails, or an entire data center has a catastrophe. Both data and log backups are completed online, with negligible impact to the users. Parameters in SAP HANA are set by default, but customers can change these to determine their own recovery point objective (RPO). Focusing on your SAP Business Suite on SAP HANA migration project, we'll cover the backup and recovery scenarios that are specific to a scale-up scenario or a single node. SAP Business Suite on SAP HANA isn't yet supported on a scale-out scenario, so we'll avoid the backup and recovery options for a multi-node environment. We'll also point out where you can enable third-party backup





options, but we won't cover those vendor tools specifically. The concepts we'll cover apply to both SAP HANA-provided tools and third-party tools.

**Note** [«]

Want to know more about available third-party certified backup options? Go to *http://global.sap.com/community/ebook/2013_09_adpd/enEN/search.html*, and enter "hana-brint" into the search box. A list of currently certified partners will be displayed.

### 3.4.1 Backups

Your organization's backup strategy will be a combination of SAP recommendations, industry best practices, and lessons learned from operating an SAP ERP system in your environment. SAP supports data backups, storage snapshots, and log backups. These backups can be used to recover the system in the event of failure to the most recent point in time or a specific point in time, or they can be used as the source for a system copy. The backup and recovery function will be performed and set up by your SAP NetWeaver (Basis) administrators in coordination with the infrastructure team that supports the backup devices and storage solutions (ideally, the SAP NetWeaver administrator attended the HA200 class we talked about earlier). SAP recommends a backup strategy that leverages all types of backups available, specifically, a daily storage snapshot, automatic log backups, and a complete data backup once a week. In the event of a failure during the week, you can either restore from the nightly storage snapshot or from the complete data backup, and then read the logs from the automatic log backup. If you have a failure in your storage solution, the storage snapshot won't be available, but the complete data backup will be because you would have saved it to a location other than the local SAP HANA system. Leveraging all backup types mitigates risk and dependencies on a single component of the backup solution.

As you dig into the administration, backup, and recovery of SAP HANA, you'll see references to the savepoint. As we mentioned, SAP HANA persists data on disk. During normal operations, changed data is automatically saved from memory to disk at regular savepoints. The time period for the savepoint is set by default to every five minutes, but this value can be changed. So, every five minutes, the savepoint is defined, and data is written to disk. If you think about a system restart, the savepoint reduces the time to restart because it doesn't have to read all redo log files. Only those redo logs that occurred after the most recent savepoint

need to be read after a system restart. It's important to know that savepoints are written asynchronously by each service within SAP HANA. A global savepoint is a consistent collection of savepoints for all services in the system. For example, a global savepoint is written when you start a complete data backup. The point here isn't to make you an expert at SAP HANA database backup and recovery, but instead to talk about critical components of the backup mechanisms. If during your SAP Business Suite on SAP HANA migration, you begin to see errors related to the savepoint, you now know this is related to the ability to recover or restore your system.

[»] **Note**

Want to become an SAP HANA backup expert? Check out the *SAP HANA Administration* book (SAP PRESS, 2014). Or see more by going to *http://help.sap.com* and choosing Technology Platform • SAP HANA Platform • System Administration and Maintenance Information • SAP HANA Administration Guide.

Now let's walk through a system backup using SAP HANA Studio:

1. Be sure to download and install SAP HANA Studio from the SAP Software Download Center at *https://support.sap.com/swdc.*Then choose Support Packages & Patches • A-Z Products • H • SAP HANA PLATFORM EDITION • SAP HANA PLATFORM EDIT. 1.0 • Entry by Component • HANA Studio.
2. Launch the SAP HANA Studio, select File • New • Folder, and give the folder a name. The folder is a logical grouping of systems. You can have one folder for the entire landscape or one folder per Software Development Lifecycle (SDLC) phase (development, QA, production). The folder only exists locally.
3. Right-click the folder you created, and click Add System. Enter the server name, instance number, and a description. Click Next.
4. Choose Authentication by Database User, and enter the credentials given to you by your SAP HANA installer. Click Finish after you've entered the user ID and password.
5. Right-click the system you added, and select Logon.
6. Right-click the Backup folder, and select Open Backup Console, as shown in Figure 3.35.





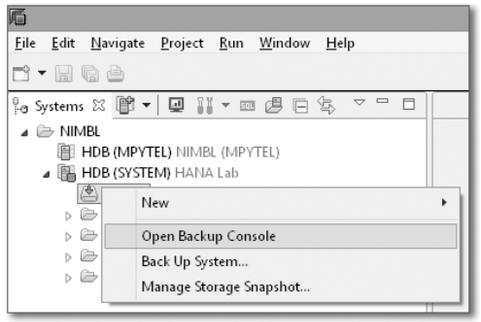


**Figure 3.35** Opening the Backup Console to View Backup History

7. The initial screen displays backups in progress, if any, and the most recent successful backup. Click the Backup Catalog tab to see the backup history detail.
8. If you want to run a backup now, then right-click the Backup folder and select Back Up System, as shown in Figure 3.36.

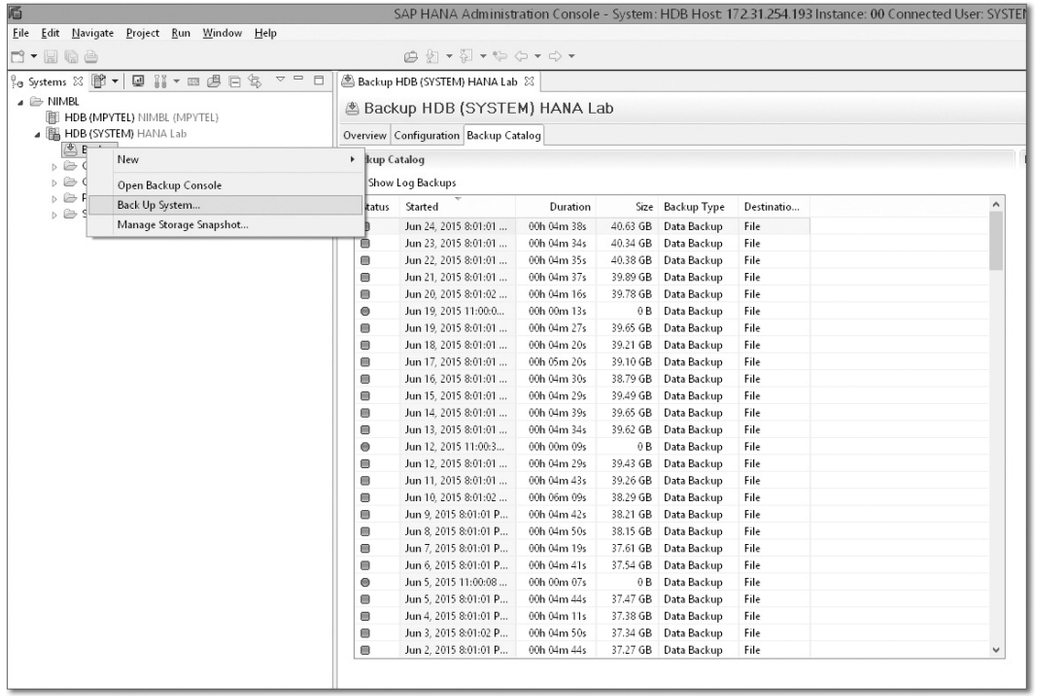


**Figure 3.36** Backup Catalog Displayed

9. Review the parameters on the SPECIFY BACKUP SETTINGS screen. Adjust the BACKUP DESTINATION or BACKUP PREFIX if required (see Figure 3.37). Click NEXT.

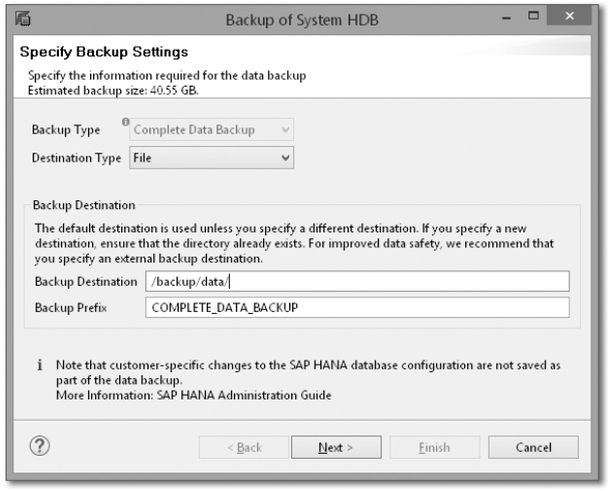


**Figure 3.37** Default Backup Location and Description

10. On the next screen, confirm the settings, and then click FINISH.
11. You get a message that the backup is running (see Figure 3.38). Wait for the status to update, and then click CLOSE.

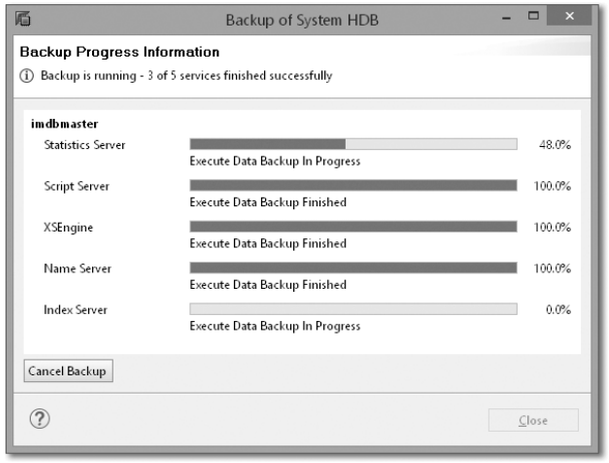


**Figure 3.38** Backup Status Displayed during the Backup Execution





12. Confirm the results, and click CLOSE.

What we've shown covers a system backup. The automatic log backup functionality will facilitate the log backups to the specified directories in the GLOBAL.INI. We've only covered the very basic backup scenario you would use during the initial phases of your project. Your backup strategy will need to be automated, which is supported by SAP HANA out of the box. All backups and restores can be run from the command line in addition to SAP HANA Studio, the Database Administration Cockpit (DBA Cockpit), or the SAP Database Control Center (DCC; more on this in Chapter 9). Thankfully, SAP has provided an excellent backup script and documentation on how to use that script with Cron in Linux. Attached to SAP Note 1651055 – Scheduling SAP HANA Database Backups in Linux, you'll find two files. One file is a compressed file with the script and configuration file. The second is a PDF document describing the usage of the backup script. The script can perform multiple functions related to backup paths, naming conventions, retention periods, backup catalog maintenance, and listing backups available. Figure 3.39 is an example from our lab environment. In this example, we keep two days of backups on a shared storage. For good measure, we move one backup to a cloud location daily. This ensures that we can restore in the event of an appliance or storage failure.

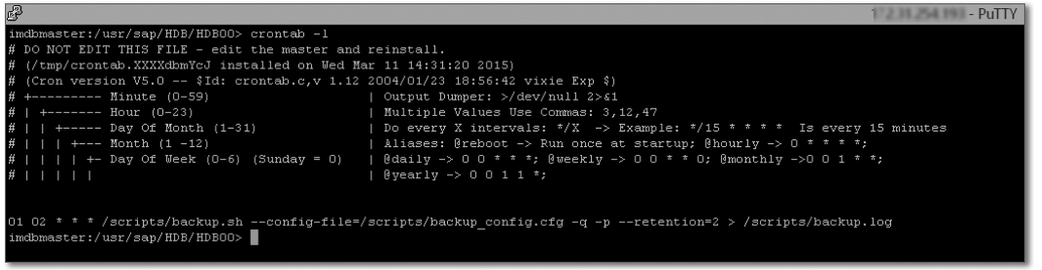


**Figure 3.39** Retaining Backups for Two Days Using the Retention Parameter

### 3.4.2 Recovery

The recovery tools in SAP HANA include features you expect from a modern database. You can recover to the latest point in time or a specific point in time, and you can use the recovery tools to restore to a new system to create a copy of the database. Unlike backups that can be run while the database is operational, the recovery process must be done while the database is shut down. Here are some other restrictions that apply when attempting to recover a database:

- At the start of the recovery, all data and log backup files must be accessible in the system or via the third-party backup tool.
- To recover the database, you need at least one full database backup.
- The recovery must be executed on a system that is the same release or higher than the source system.
- You can't pause and resume a restore. After the restore starts, you can cancel; however, the database will be left in an inconsistent state.
- There are no restrictions on restores from multinode systems to single-node systems or the reverse.
- The permanent license is restored if the system ID and landscape ID haven't changed. If either has changed, a temporary license will be installed that is valid for 90 days.

In Figure 3.40, you can see three options to recover your SAP HANA database. During the SAP Business Suite on SAP HANA project, you may want to restore the database in multiple scenarios.

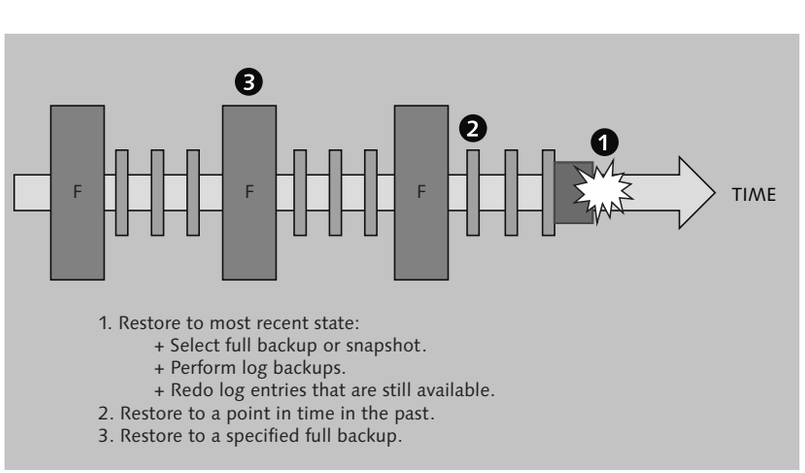


**Figure 3.40** Example Restore Scenarios

In one example, you might want to run your cutover procedures multiple times to develop the shortest execution strategy. In this scenario, you complete an SAP Business Suite on SAP HANA migration, take a full backup, and then let your team execute the cutover plan. After they finish, you can validate the system. If the timing was too long, or something wasn't correct, you'll need to restore the database





and try again. Without the backup and restore scenario, you would have to run the entire migration again. Another scenario includes a QA refresh after the production go-live. In this scenario, you execute your go-live cutover procedures and system validation, and right before you release the system to the end users, you take a backup. You then restore that backup to the QA system as a system copy to provide a working copy of production to support defect resolution.

Your SAP NetWeaver administrators will become very familiar with your backup and restore process. Now, let's walk through an example database restore using SAP HANA Studio. There are command-line options to perform the restore as well.

1. Log on to the SAP HANA database using SAP HANA Studio, which you downloaded and installed in Section 3.4.1.
2. Right-click the system, and select Backup and Recovery • Recover System (see Figure 3.41).

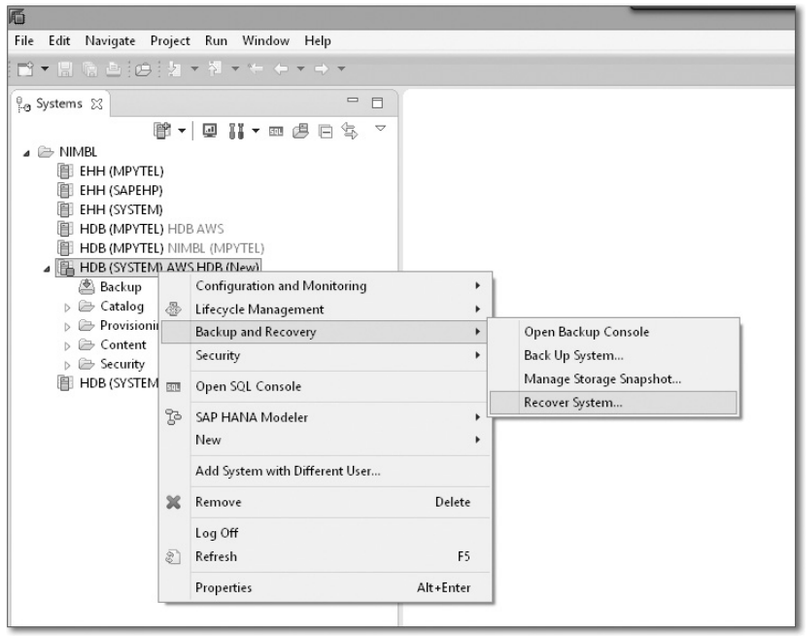


**Figure 3.41** Starting a Database Restore on SAP HANA

3. Enter the <SID>ADM user ID and password if prompted.
4. Confirm that the database can be shut down.

5. Specify the recovery type by selecting Recover the database to the following point in time to simulate the restore you might run during the project (see Figure 3.42).

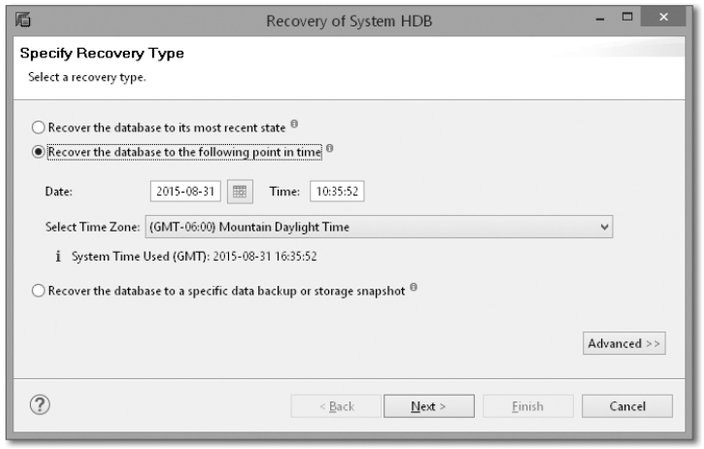


**Figure 3.42** Selecting the Point in Time for Recovery

6. Click Next, and choose the Locations of the backups (see Figure 3.43). You can add multiple locations here, for example, the "cold" storage location.

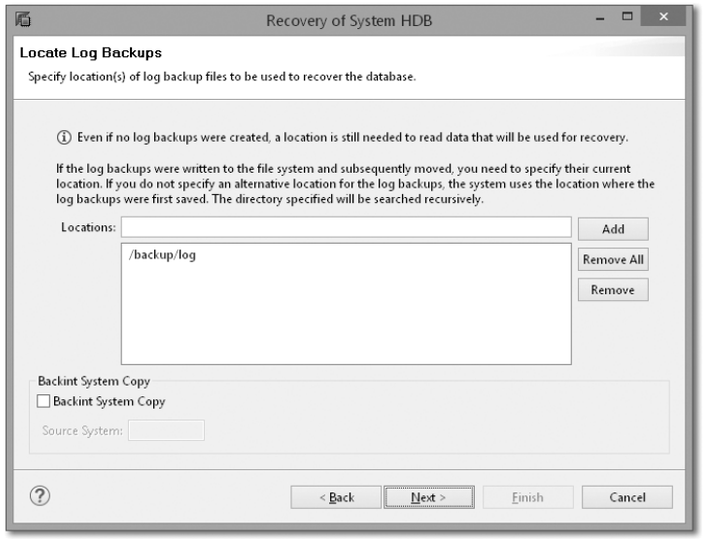


**Figure 3.43** Selecting Locations for Backups





7. An overview of the data backups is displayed (see Figure 3.44). You can manually input an Alternative Location, and then click Check Availability.

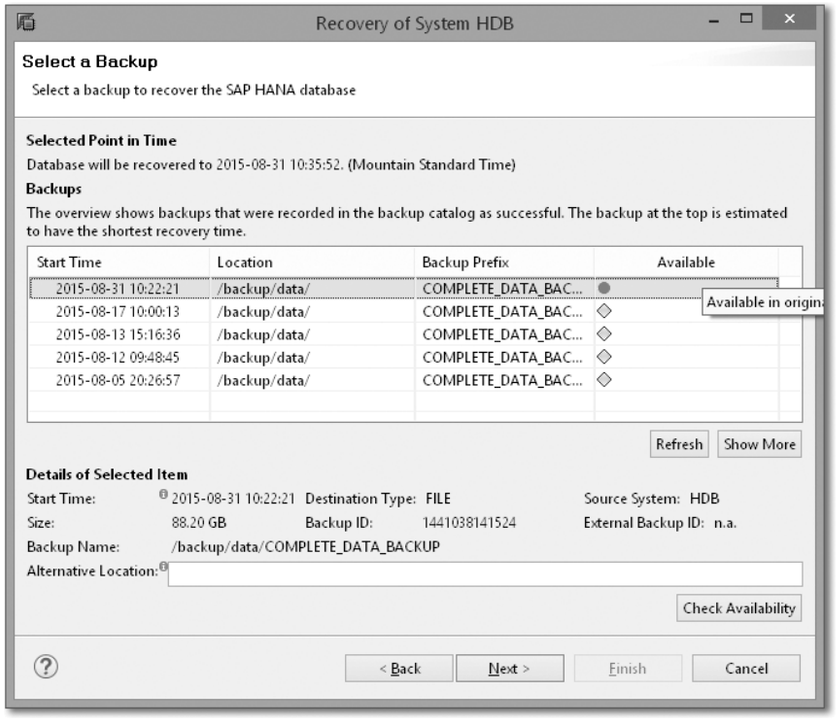


**Figure 3.44** Confirmation that the Backup Is Available

8. Select the data backup, and choose Next.
9. The wizard now lets you choose some additional options. You can select Use Delta Backups to potentially shorten the restore time if they are available.
10. Click Next after you've confirmed your options (see Figure 3.45).


**Figure 3.45** Additional Restore Options for SAP HANA

11. Review the summary screen, and click FINISH to begin the restore. You should see a progress screen like the one shown in Figure 3.46.
12. If successful, you'll see the message SYSTEM <SID> RECOVERED (see Figure 3.47). Click the CLOSE button, and confirm that SAP HANA services have been started.





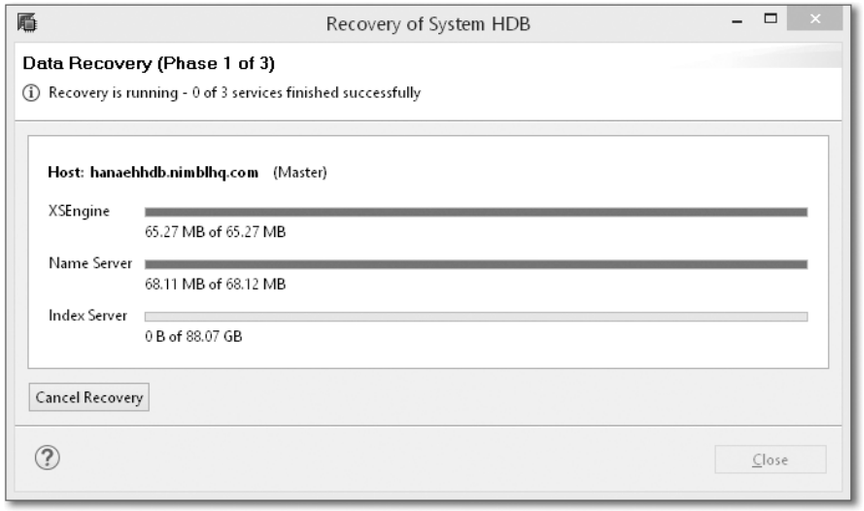


**Figure 3.46** Restore Status Screen

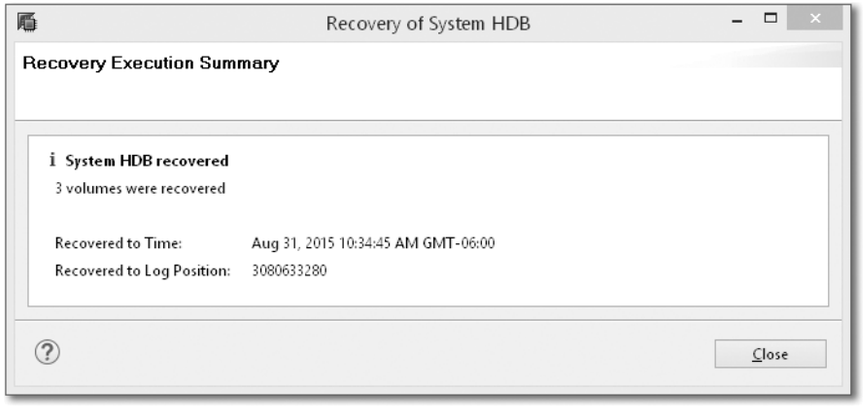


**Figure 3.47** Summary Screen Confirming Restore Complete

## 3.5 Summary

In this chapter, we focused on some of the infrastructure requirements needed to prepare your landscape for the SAP Business Suite on SAP HANA migration. Your organization will need to manage multiple types of change during the project.

Your technical team will be managing a new hardware platform, learning to manage a new database, and potentially making changes to a landscape that has been stable for several months or years. These will be challenging times. However, you can mitigate risk by following the guidelines and best practices published by SAP. Performing the ABAP and Java stack split early on, in addition to the enqueue server split, will save time and reduce risk during the actual migration to SAP Business Suite on SAP HANA.

In the sizing section, you learned about the varied sizing activities you can perform. You can initially perform a rough sizing; however, the ABAP report provided by SAP is so simple to install there really isn't a need to guess. It estimates the required size of the SAP HANA appliance from your current production environment, to which you can add the SAP recommended 20% for safety purposes. Finally, in the sizing section, we talked about the user-based and throughput-based sizing. The user-based sizing requires more time than the ABAP report because you need to sort your users by functional area into activity types. The more accurate throughput-based sizing requires significantly more effort due to the extensive data you need to collect about your organization. Remember that bad sizing estimates can be mitigated by running a performance test during your project. Performance testing tools can now be purchased on a term basis (e.g., 60 days), utilized for the project, and then decommissioned after your project is complete. The performance test is the insurance policy for not collecting or focusing enough energy on a throughput sizing.

We ended the infrastructure topic on backup and recovery. While this section was merely an introduction to the options available with SAP HANA SP 09, you should now appreciate how important those tasks will be to your project. Depending on your landscape strategy (five-system or four-system), your team might be performing several restores or system copies over the life of the project. The features we covered in this chapter were focused on SAP HANA SP 09 and below. We know that SAP has planned innovations in SP 10 to include new backup types and capabilities, but your SAP Business Suite on SAP HANA backup and recovery strategy probably won't need to change unless the business is asking for higher Service Level Agreements (SLAs). You still have the ability to perform periodic weekly, monthly, and ad hoc full backups whenever required. The policy governing when and why SAP ERP is backed up won't change. However, the procedures and documentation of the tools will need updating.

# Contents

**PART IV Optimizations**

# Index

## D



## E



## F



## G



## H



## I



## J



## K



## L



## M

## N



## O



## P



## Q



## R



## S

## T

## U



## V







## W



## X

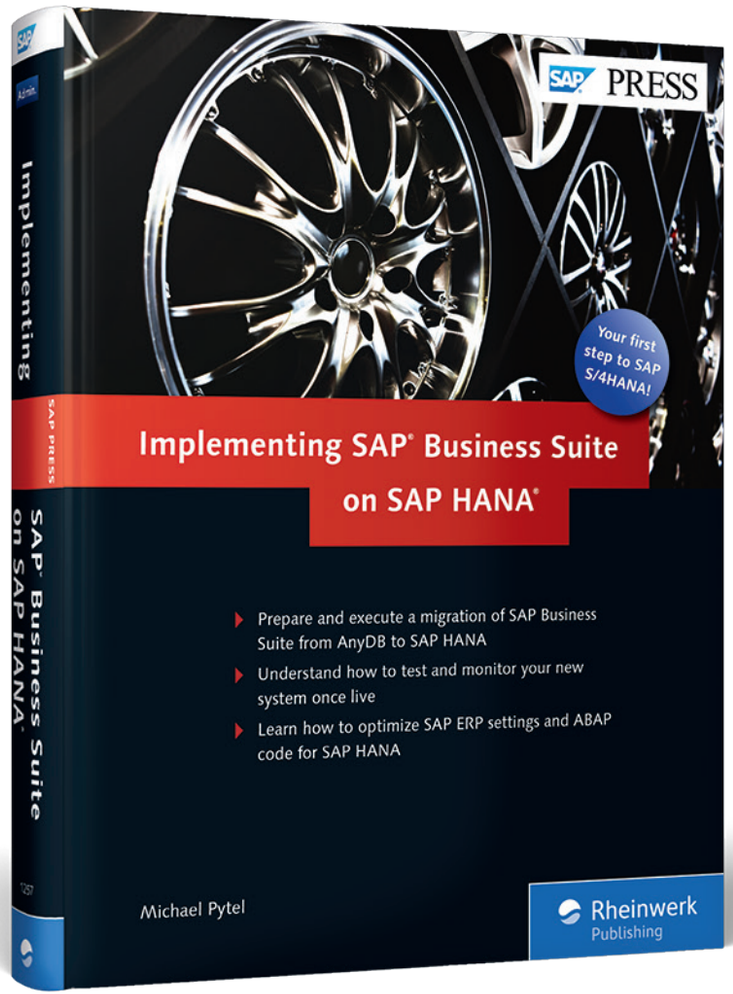


**Michael Pytel** is currently the chief innovation officer and co-founder of NIMBL, an SAP implementations consulting firm. He has more than 14 years of hands-on technical experience with SAP solutions, including SAP HANA.

Michael Pytel

# Implementing SAP Business Suite on SAP HANA

597 Pages, 2016, $79.95/€79.95
ISBN 978-1-4932-1257-6

www.sap-press.com/3895

*We hope you have enjoyed this reading sample. You may recommend or pass it on to others, but only in its entirety, including all pages. This reading sample and all its parts are protected by copyright law. All usage and exploitation rights are reserved by the author and the publisher.*