Improve the flexibility of your SAP solution landscape with the new Adaptive Computing Controller 7.1

by Nils Krugmann

Nils Krugmann
Technology Consultant,
SAP Consulting
Integration Solutions
Management,
SAP Germany

Nils Krugmann is a technology consultant for SAP Consulting. He graduated from the University of Tübingen, Baden-Württemberg, Germany, with studies in computer science and computational biology. He joined SAP Consulting in 2006 with a focus on SAP NetWeaver, ERP, infrastructure, and SAP Basis technologies. Adaptive computing became one of his major consulting focuses. Currently, he is covering the Ramp-Up of ACC 7.1 from the consulting side. You may reach him at nils.krugmann@sap.com.

The demands on IT are growing exponentially while everyone expects technology to work silently. When it comes to SAP systems, your users expect you to react to changing requirements as soon as they have clicked on the Send button in their email clients to ask for them. When you introduce new applications into a system landscape, you tend to heavily oversize your hardware because future performance demands often remain uncertain. However, when your applications and their demands increase, you typically move your instances to higher-performance hardware that might already exist in your landscape or that may need to be added.

The new Adaptive Computing Controller (ACC) in SAP NetWeaver 7.1 can help you address the needs of your users quickly, easily, and seamlessly while you manage and optimize performance. The ACC is the right tool to react to foreseen and unforeseen changes in your system landscape while hiding the complexity behind it. The new ACC is the central point of control from which you can manage an adaptive computing landscape.

Note!

The Ramp-Up version of ACC 7.1 is running, and the software is expected to ship with the SAP NetWeaver Process Integration (SAP NetWeaver PI) 7.1 installation media. It's not part of SAP NetWeaver PI; it's just in the same shipment. The planned release date for the general availability of ACC 7.1 is June 2008. When installing it from the SAP NetWeaver PI installation master, an SAP NetWeaver Java 7.1 system with deployed ACC will be installed.

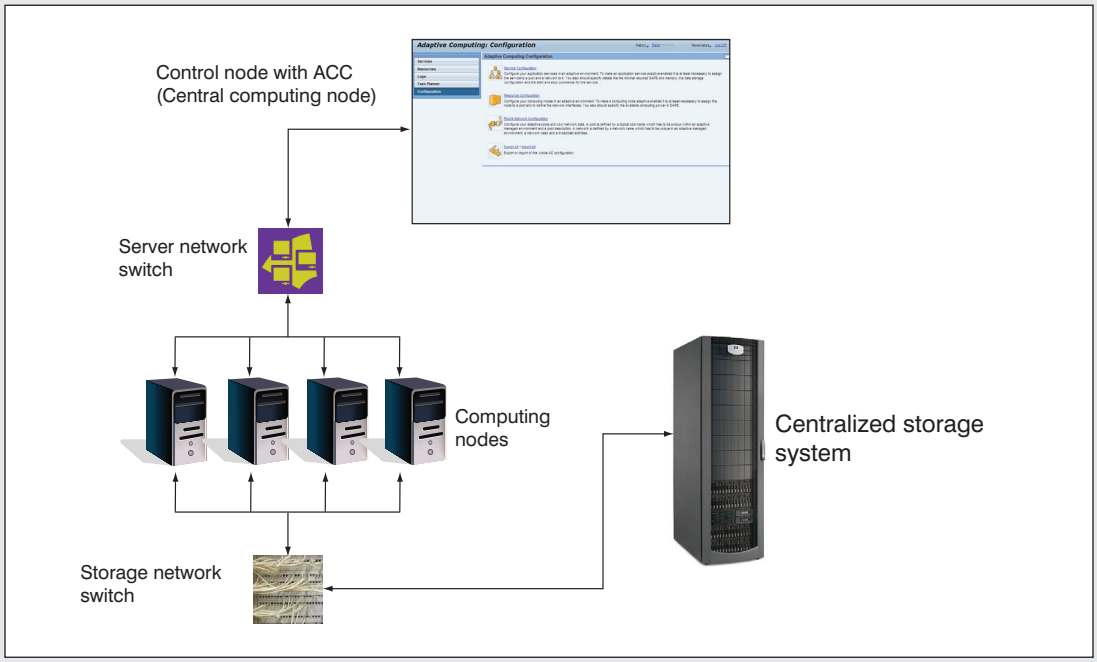

Figure 1 An adaptive computing landscape

Adaptive computing is an approach to dynamically assign hardware resources to your SAP application instances. If a particular server doesn't fulfill the performance requirements of the application you plan to run on it, the ACC enables you to quickly move the instance to more powerful hardware. From now on, you don't need to scale your servers for potential upcoming demands several years in the future. With ACC 7.1 you can easily replace outdated hardware by moving your instances to higher-performance replacement servers that fit the growing demands of your applications.

This article shows you how adaptive computing can help you to improve the flexibility of your SAP system landscape. The article begins with an overview of adaptive computing. You will learn about the adaptive computing landscape, the meaning of virtualization and its role in adaptive computing, the ACC upgrade process, and so on. Then, the article walks you through the tasks required to install ACC 7.1 on SAP NetWeaver, including configuring the ACC instance, reporting instances to the System Landscape Directory (SLD), and configuring the application services. Finally, you'll learn how to securely move a system with the ACC.

Let's begin by exploring what adaptive computing is and the benefits to be gained with implementation.

Adaptive computing overview

In an adaptive computing landscape, you can quickly and easily relocate applications from the server they are running on to any free server in your system landscape. The technique used for this relocation process is virtualization, which enables instances to move from one server to another.

As shown in **Figure 1**, the server network switch connects the computing nodes to each other. With the ACC, you cannot move instances beyond network boundaries so you must configure your computing nodes on the same network. Since the ACC node must

How to jump-start adaptive computing

The shift from a traditionally managed SAP system landscape to an adaptive computing managed landscape can require implementation and configuration efforts. You need to virtualize the instances running on physical host names, enter all your computing nodes into the ACC, and prepare for central OS deployment. Perhaps you want to start small, become experienced with the concept, and then expand it slowly throughout your system landscape. In my experience as a technology consultant and through numerous ACC implementation projects, I recommend the following:

1. Integrate your existing dialog instances (DIs) in the ACC by installing new ones on virtual host names for your adaptive computing landscape.
2. Move the application data and the components onto a central storage system, and integrate the new DIs into the ACC.
3. Use logon groups to push your users to the new instances. After you have relocated all your users to the new instances, you can delete the old ones. This is a small move and presents no danger of an interruption in operations.
4. After you become accustomed to configuring and moving application servers and managing them with the ACC, start to virtualize the central instances (CIs) of the new installations and attach them to the ACC.
5. Move existing CIs to virtual host names.
6. Integrate the database (DB) instances.

You can now fully benefit from the adaptive computing concept. In the end, you may want to introduce multiple instancing as well (i.e., running more than one SAP instance on a computing node), but that discussion goes beyond the scope of this article.

have access to your nodes (it connects via a Web service), it also needs physical access to that network.

Depending on the storage technology you're using, you may need a storage network switch to connect the nodes to central storage. Fiber-channel and Internet SCSI (iSCSI) controllers in your nodes must somehow connect to central storage, and you can achieve that by using a storage switch. If a Network File System (NFS) is used, you don't need a storage network switch because NFS uses the IP network to attach to the exported storage locations. To find out if your storage solution is supported during the Ramp-Up phase, please ask your technology partner since there might be some restrictions. ACC 1.0 supported fiber-channel- and iSCSI-based solutions. (If you are new to adaptive computing, see the sidebar on this page for my recommendations on how to approach it. If adaptive computing is already part of your landscape, see the sidebar on the next page to find out about the benefits of ACC 7.1 and the differences between using ACC 1.0 and using ACC 7.1.)

Note!

Starting with SAP NetWeaver 7.1, SAP is changing the naming conventions for CI and DI. CI will be called "primary application server instance"; DI will be called "additional application server instance."

ACC 7.1 vs. ACC 1.0

If you're already familiar with the adaptive computing concept, you might wonder how you can benefit from the new ACC in SAP NetWeaver 7.1 and the differences between ACC 7.1 and ACC 1.0.

- For ACC 7.1 you no longer need a Secure Shell (SSH). All communications between the new ACC and the computing nodes will be made through real Web services (see SAP Note 927637 – Web service authentication in sapstartsrv as of Release 7.0 for details).
- You no longer need SAP Solution Manager to enable application services for adaptive computing. Simply set up your nodes and instances within the ACC itself. SAP has completely rewritten the host agents, and they have new maintenance and configuration functions, such as an automatic update if the currently running version is outdated.
- You are now able to manage Java central services and ABAP central services as you do your own SAP instances. SAP has improved the usability of the central services and will extend these functions into new functionality (such as management functions for virtual machines) in the near future.
- In ACC 7.1 SAP has improved the monitoring of service availability so it will now be checked on the database level and through direct access to the SAP NetWeaver Administrator (NWA) content.
- You can monitor, start, and stop instances that are running on a physical host name but have not yet been virtualized by executing them on an alias host name with ACC 7.1.
- ACC can also manage standalone, enqueue replication servers, beginning with Service Pack (SP) 5 of ACC 7.1.

Instead of using the actual host name of a specific underlying server when you install SAP software, you can now use a virtual host name in adaptive computing. You map the virtual host name to a virtual IP address on the domain name server (DNS; actually, the IP address is virtual, not the host name itself). After that, you can activate (or deactivate) an IP address on the network interface (NI) belonging to the server on which you want to run the instance. Virtualizing instances (i.e., making them virtual) frees them from any dependency on the physical servers on which they were originally installed. It also gives you great flexibility and opportunity to extend and manage the performance of your entire SAP system landscape (from SAP R/3 4.6C through SAP NetWeaver 7.1, as well as SAP Business ByDesign systems, APO/LiveCache, and TREX instances).

The ACC manages virtual host names and allocates storage resources for the servers enabled by adaptive computing. The ACC is also responsible for starting and stopping application instances. In this context, an application instance can be a CI, a DB instance, or a DI. The ACC "hides" the complexity involved in relocating instances. The following are some of the benefits your company can gain from adaptive computing:

- **Server pooling:** After you have set up your servers to operate in an adaptive computing landscape, you have one flexible pool for all your SAP systems. You can then start your systems on the server with the most available resources, the one that is best suited to the instance at the time.
- **Idle systems:** You can stop idle systems (e.g., data-archiving systems, legacy systems) and wake them up when you need them. Those systems won't be allowed to needlessly allocate servers. For example, you can:

  - Stop a system and later restart it on any free server when the system is needed again.
  - Use free servers to start additional application instances for other systems or as a resource that lets you start additional training or test systems.
  - Run test systems on available hardware, and stop development and training systems overnight.
  - Employ free servers for DIs that support nightly batch operations.

- **Easy and fast hardware replacement:** If you want to upgrade a server with more memory or additional CPUs or exchange broken hardware, you can simply: relocate all application services running on that server to a free server within your landscape; upgrade or exchange the server with a new one; prepare the server for adaptive operation; and move the formerly running services back to the original server.
- **Hardware testing:** If you plan to introduce new hardware, you can add test hardware to your landscape and relocate an instance to see how it will behave on the new server.
- **Mass operations:** This feature lets you stop and start all the instances in your system landscape in parallel. In large system landscapes with dozens of instances, this functionality is useful when you have to shut down all your systems for hardware maintenance (say, on Friday evening) and then bring them all back up again (say, on Sunday evening).
- **Training systems:** You can switch training systems off in the evenings and on again on the weekends. That way, you can use your hardware resources for your production systems at times when they are heavily used.
- **Enhanced availability:** Relocating an application instance still includes the downtime required to shut down the instance, activate virtual host names and storage resources on the new server, and start the instance there. However, you will find that the new ACC can bring enhanced availability to your system landscape.
- **Performance peaks:** Because you typically need to use your systems heavily during month-end or fiscal-year closing, and not evenly throughout the month, you can now automatically move those heavily used instances to faster, better hardware during their peak times and then back to slower, lesser hardware for the rest of the month.
- **Time-critical operations:** For operations such as upgrades, Unicode conversions, or SAP Test Data Migration Server (TDMS) runs, you can move the instance to more powerful hardware and then move it back to a system that best suits its typical demands after the critical operation is complete. You can use the ACC to address these critical operations.

ACC upgrade process

The new ACC 7.1 can manage SAP NetWeaver 7.0 systems, but it cannot run on such a system. The ACC application is loosely coupled in the SAP NetWeaver Administrator (NWA) in SAP NetWeaver 7.1, so it needs an SAP NetWeaver 7.1 system to run it. And since the new ACC is integrated with the NWA, you can't deploy or install the ACC part of the software separately. Before you can install the ACC, you need to make sure that the host system has been upgraded.

To get ACC 7.1 running, you need to install it from the SAP NetWeaver PI 7.1 installation media. If you already use ACC 1.0, which runs on SAP NetWeaver Application Server Java (SAP NetWeaver AS Java) 6.40, you cannot upgrade the ACC from version 1.0 to version 7.1 directly; you need to install a whole new system, which consists of SAP NetWeaver AS Java 7.1 and the ACC application integrated into the NWA. The upgrade process would look like the following:

1. Install the ACC 7.1 out-of-the-box (no direct upgrade from ACC 1.0 is possible).
2. Export the SLD content from the SLD that was attached to ACC 1.0. Then, import the SLD content from SLD 7.1, which is necessary to run ACC 7.1. You cannot use any older version of SLD, so you must either upgrade SLD 6.4 or 7.0

to SLD 7.1 (which isn't possible at present, but should be by the end of 2008) or set up data-forwarding from the old SLD to the new SLD.

3. Run the script to update the SLD content after importing the data (as described in the ACC upgrade documentation on the SAP Help Portal at http://www.help.sap.com and on SDN at https://www.sdn.sap.com/irj/sdn/adaptive → Adaptive Computing Controller (in SAP NetWeaver 7.1) → Under help topics, click on Adaptive Computing Controller (in SAP NetWeaver 7.1) → Upgrading the Adaptive Computing Controller).
4. Upgrade the SAP host agents to the new version of SLD, and configure them to report to the new SLD.
5. Provide an instance agent for the managed instances of releases older than SAP NetWeaver 7.0 (see Adaptive Collective Notes for details; 725397 – Adaptive Computing for ACC 1.0 and 1008828 – ACC 7.1 PI/Adaptive Computing Controller Collective Note). The 6.40 version of the instance agent will apply to older kernel releases.

At first, you install ACC 7.1 on the central computing node (**Figure 1**) and configure it with a local SLD, which is already part of your SAP NetWeaver installation. If your SLD is version 7.0, you must either set up bridge-forwarding between the old and new SLDs or upgrade your current SLD if you don't want to use the new SLD that comes with ACC 7.1. For example, you might be using your old SLD for applications such as SAP NetWeaver PI, SAP NetWeaver Development Infrastructure (NWDI), Web Dynpro (adaptive Remote Function Calls — RFCs), etc. Because ACC 7.1 is in Ramp-Up, I don't recommend you use SLD 7.1 for other productive applications at this time. You also need to install an SAP gateway for all the RFC-based communications in your landscape.

Another ACC requirement is to use a central storage system. In general, the ACC supports network-attached storage (NAS), iSCSI, and fiber-channel-based storage systems. For iSCSI and fiber-channel-based storage, you need a special partner library based on your storage solution. (You can find out how to obtain the library from your hardware vendor from SDN at https://www.sdn.sap.com/irj/sdn/adaptive.)

The ACC also requires that you install your application services on virtual host names. The example in **Figure 2** shows you how to manually activate a virtual IP address under Linux on NI `eth0`.

In this code snippet, after the virtual IP has been activated on `eth0` with the `ifconfig ... up` command, you can see the result of this activation by executing `ifconfig` only (without any parameters). This shows you all the available NIs in your system, as well as any virtual IPs activated on those NIs. `eth0:10` is the virtual IP that was initially activated on `eth0` and is now active there.

Then, you have to map the virtual IP address to the virtual host name on which your instance runs by entering it into your DNS or just by editing and synchronizing the local copies of `/etc/hosts` (an IP address mapped to an alias host name). The virtual IP for the CI is 172.23.221.110; I assigned the alias host name "cidm1" to that IP:

```
acvml2211:~ # cat /etc/hosts | grep
          172.23.221.110
cidm1
```

Note!

Bridge-forwarding means setting up a data bridge between two SLDs. On the sender, you define the data bridge and then all data that the sender SLD receives is automatically pushed to the receiver SLD. Since you cannot (at the time of this publication) upgrade an SLD to version 7.1, although ACC 7.1 requires it as part of the ACC installation, you might want to keep your systems attached to the old SLD. Then, you set up a data bridge from SLD 6.4 or 7.0 to the new SLD 7.1 on the ACC system.


```
ifconfig eth0:10 172.23.221.110 netmask 255.255.0.0 up

ifconfig
eth0:10   Link encap:Ethernet  HWaddr 00:0C:29:9B:4A:1F
          inet addr:172.23.221.110  Bcast:172.23.255.255  Mask:255.255.0.0
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          Base address:0x1070 Memory:f4820000-f4840000
```

Figure 2 Activating a virtual IP address

This code snippet assigns the virtual IP address to host name cidm1. When you install an application service, you can use the command `./sapinst SAPINST_USE_HOSTNAME=cidm1` to install the CI on the server currently indicated by the virtual host name.

When you virtualize application instances, you can take the physical host name of the application instance and convert it to a virtual one (by assigning a virtual IP address) or change the host name of the instance manually or by using a system copy. (For more information on how to change the host names of existing instances, look at the corresponding guides available from SAP Service Marketplace and relevant SAP Notes in the sidebar on this page. SP 5 guides are available now; however, in the future, there might be no more printed guides (PDFs), since it is difficult to keep them updated. All documentation is available on the SAP Help Portal at http://help.sap.com/.)

You also need to have an SAP gateway running in your application landscape through which host agents running on the computing nodes can register in the local SLD of the ACC as well as the ABAP instances themselves via transaction RZ70 (SLD Administration). In addition, you have to register the application instances in the local SLD. After making these preparations, you must enable the computing nodes and instances in the ACC, so you can relocate an instance between two configured servers.

References on how to change host names

- SAP System Copy: http://service.sap.com/systemcopy
- SAP Systems in Switchover Environments: http://service.sap.com/ha
- Network Setup: http://service.sap.com/network
- SAP Note 8307 – Changing host name on R/3 host: What do you do?
- SAP Note 574310 – iSeries: specific actions when renaming DB host
- SAP Note 757692 – Changing the host name for J2EE Engine 6.4/7.0 installation
- SAP Note 403708 – Changing an IP address

Note that you need logon credentials for SAP Service Marketplace.

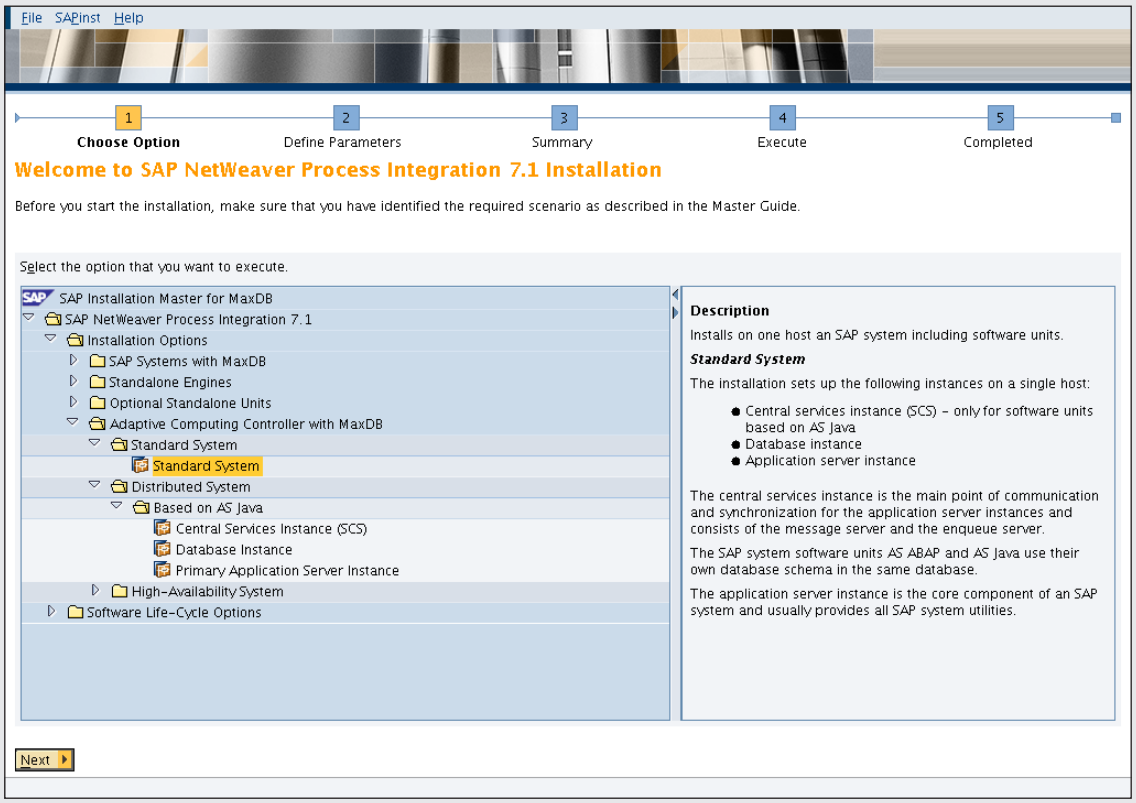

Figure 3 Locating the installer for adaptive computing

Installing the ACC in SAP NetWeaver 7.1

For the purpose of this article, let's use a system that runs Linux and SAP MaxDB as a base (Microsoft Windows isn't supported for SAP Ramp-Up). You can find the installer for ACC 7.1 (as shown in **Figure 3**) under SAP NetWeaver PI 7.1 → Installation Options → Adaptive Computing Controller with MaxDB → Standard System in the installation tree. (You need to download the SAP NetWeaver PI 7.1 installation master to install ACC 7.1.) This system is a standard system in which all the components reside on one host. (You can find more details regarding the setup of an SAP NetWeaver PI 7.1 system in the SAP installation guides at https://service.sap.com/installNWPI71 → 2 – Installation – SAP NetWeaver Systems → Installation Guides – Adaptive Computing Controller 7.1.)

To install the new ACC 7.1, you need to perform the following tasks:

- Configure the ACC instance
- Report instances to the SLD
- Add a managed system to the NWA
- Install the SAP host agents
- Add instance and servers to the ACC
- Configure the computing pools and networks
- Add computing resources to the adaptive computing pool

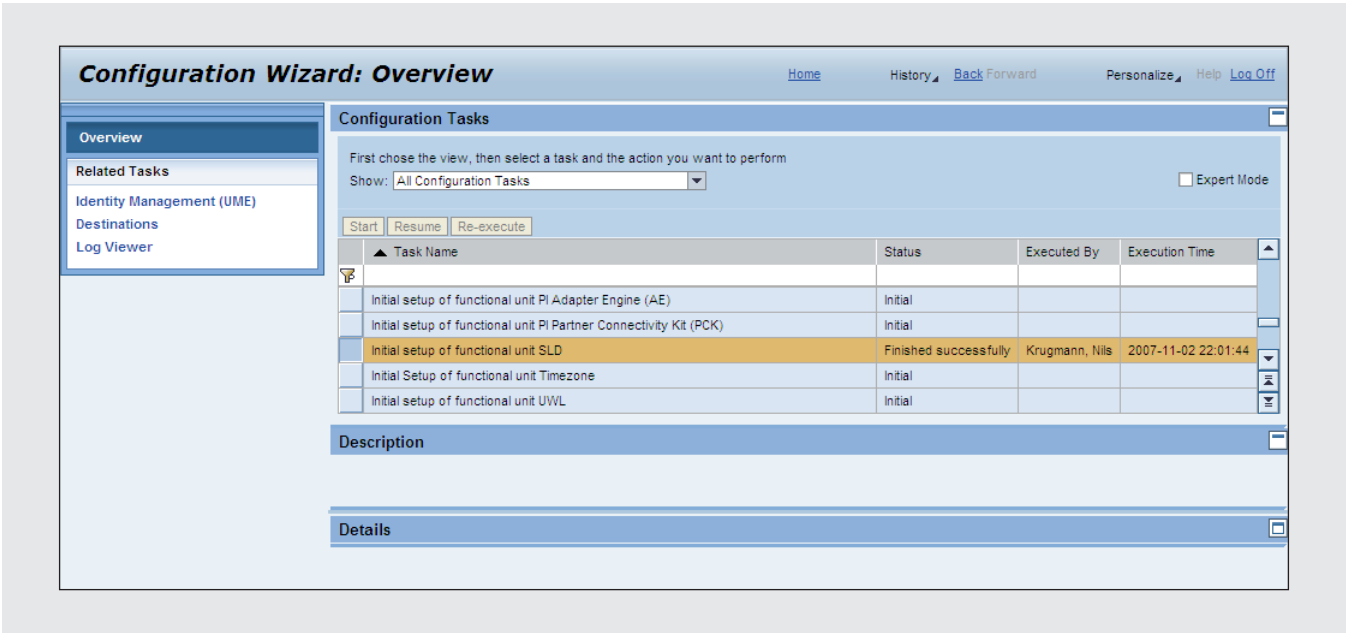

Figure 4 Configuring a working SLD for your system

Configure the ACC instance

With the help of the NWA, you can configure the ACC instance according to its future usage. Since you need SLD 7.1 for adaptive computing, you must configure the local SLD for the ACC instance. You can use the Configuration Wizard to do this. (I don't show every step of the wizard, just the options you need to select. The wizard will guide you through the details of each option.)

You call the NWA of your ACC system using the link http://<*myhost*>:5*nn*00/nwa. Then, go to Configuration Management → Scenarios. The Configuration Wizard is the new name of the Template Configuration Tool used in ACC 1.0. On the initial Configuration Wizard screen, click on Configuration Wizard, and then click on the Initial setup of functional unit SLD task (see **Figure 4**), which will highlight in orange and the appropriate button Start, Resume, or Re-execute will become active. Click on the appropriate button to continue.

- **Start** will start a wizard for the first time.
- **Resume** will reset the wizard to the point where you canceled its execution.
- **Re-execute** will run through the wizard again. (Since SLD was configured initially, you cannot re-execute this wizard after the task has completed successfully.)

On the next screen, click on the Configure NWA for Java System task (see **Figure 5** on the next page) and click on the Start button if you're running the

Note!

After the initial SLD configuration, don't forget to update the Common Information Model (CIM) by downloading the actual SLD 7.1 CIM content from the SAP Service Marketplace at http://service.sap.com/. Because the ACC makes heavy use of the SLD's functionalities, you should always import the actual CIM content in your SLD. For more information on updating the SLD models, see SAP Note 669669 – Updating the SAP Component Repository.

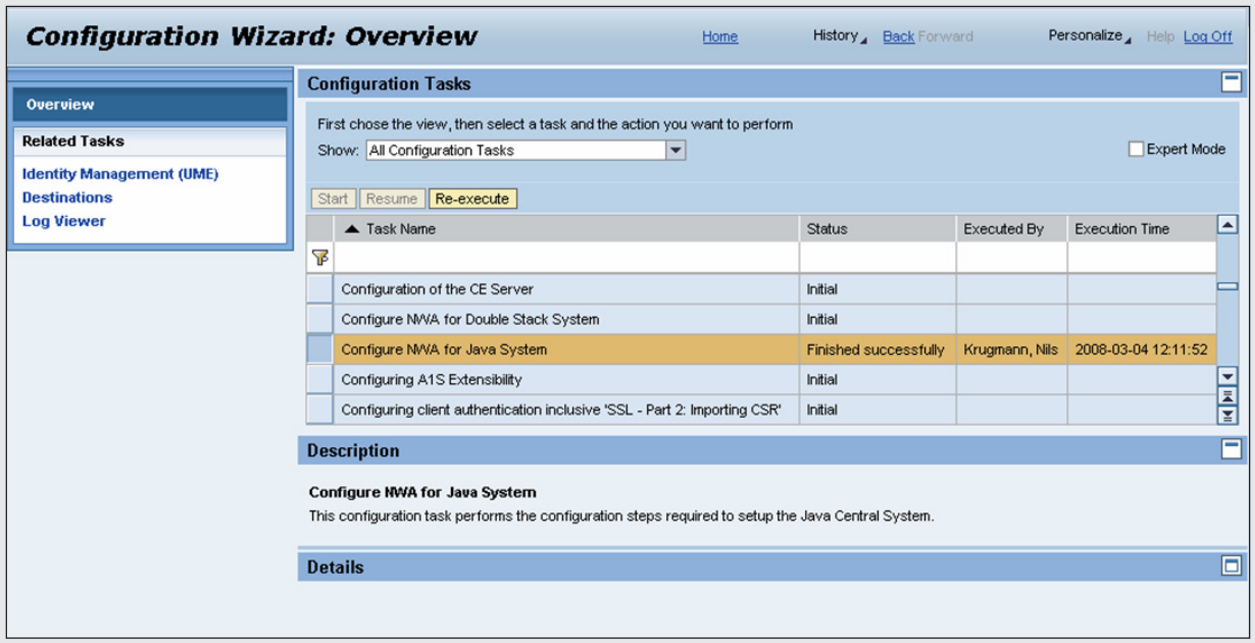

Figure 5 Setting up a working NWA for your system

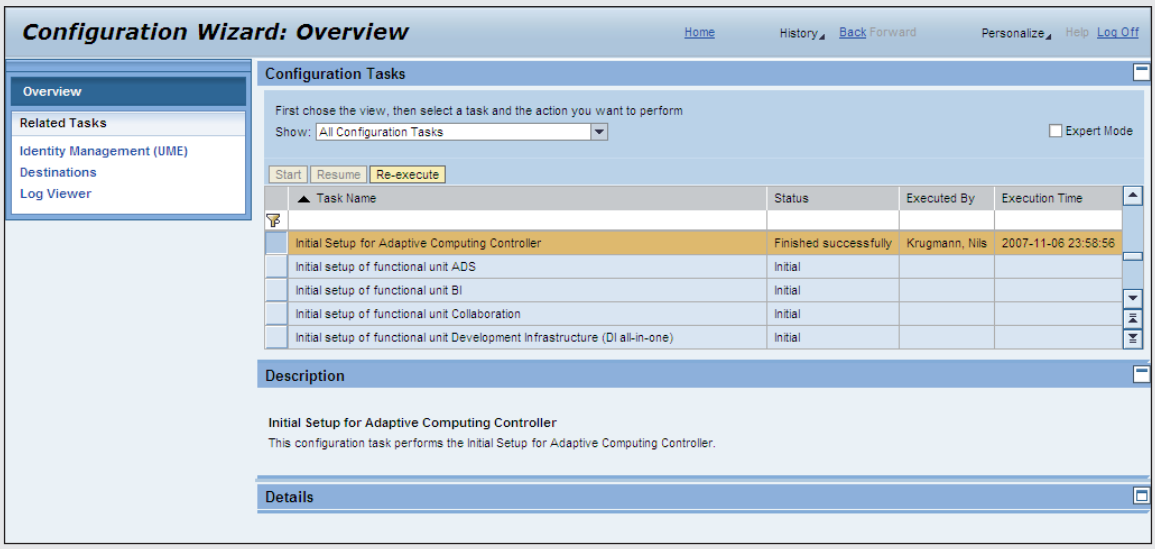

Figure 6 Setting up the ACC

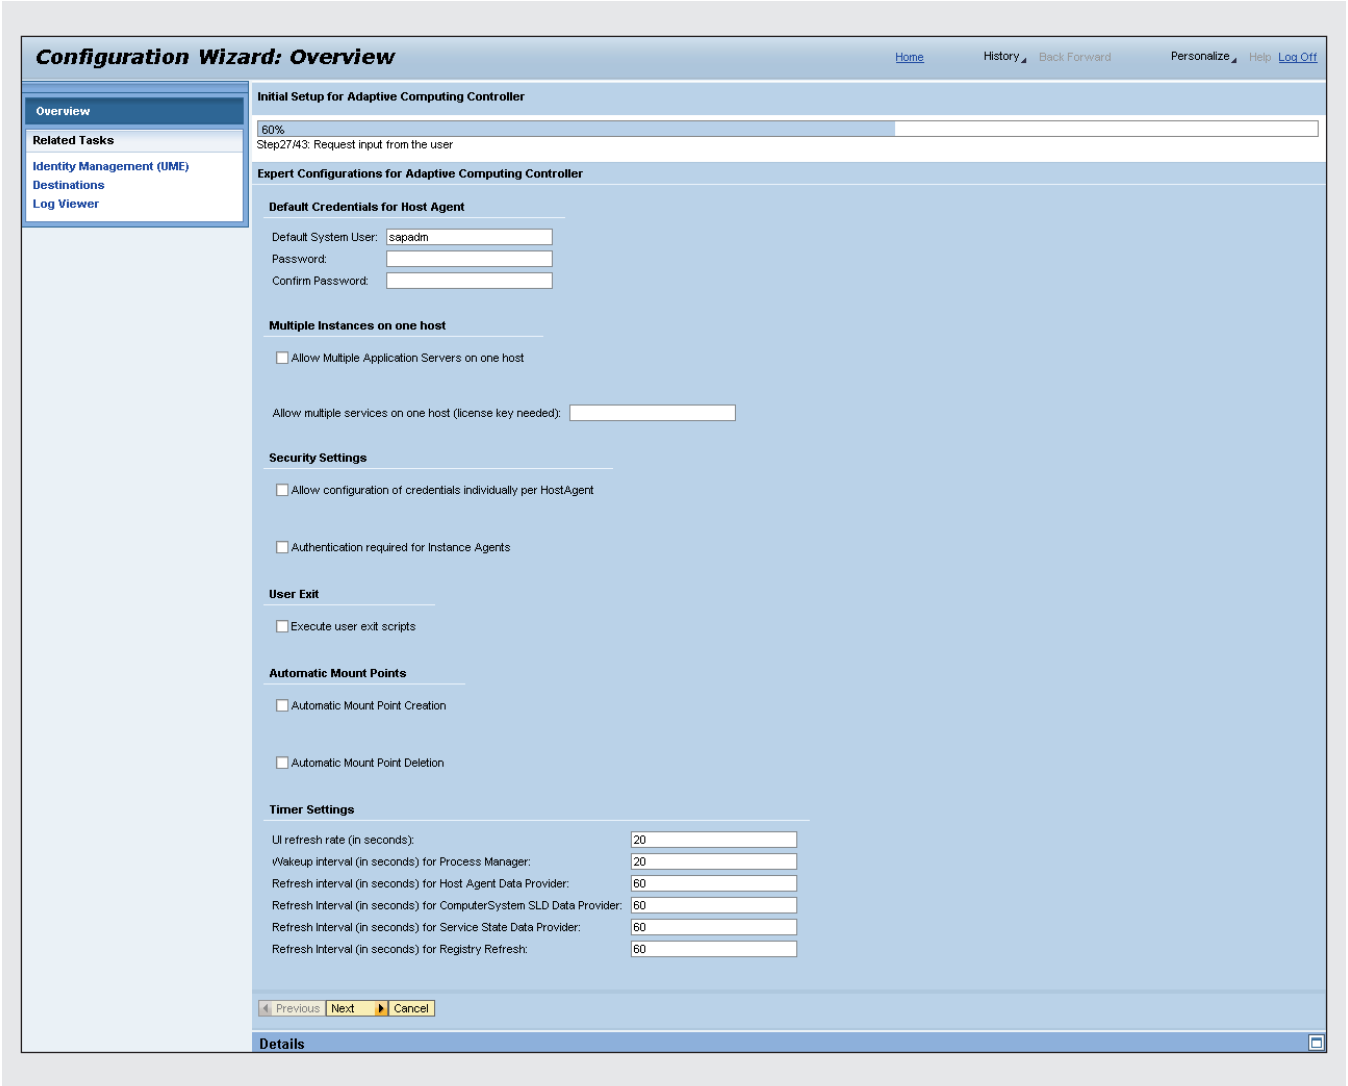

Figure 7 Setting ACC options in custom mode

wizard for the first time, or the Re-execute button if you already executed the wizard successfully and you run it again afterwards. If you configure the NWA, you define an admin landscape to which the systems will later be added (in the wizard where you add a managed system to the NWA, it automatically assigns the current admin landscape, as defined here). You change the admin landscape's name by re-executing it. The wizard guides you through the configuration steps and sets up a working NWA for your system. This step isn't mentioned in the official guides, but it's necessary.

Now you need to set up the ACC. Click on the Initial Setup for Adaptive Computing Controller task (see **Figure 6**). Then click on Re-execute.

The wizard activates ACC functionality and asks for initial parameters that you can set to the proposed standard values from SAP in standard mode. If you run in expert mode, you'll need to use your own values. Examples of those values are the GUI refresh time and the option to switch on user exits, which are deactivated by default. The screen in **Figure 7** appears when you execute this wizard in custom mode. You can then set the following options by clicking on the

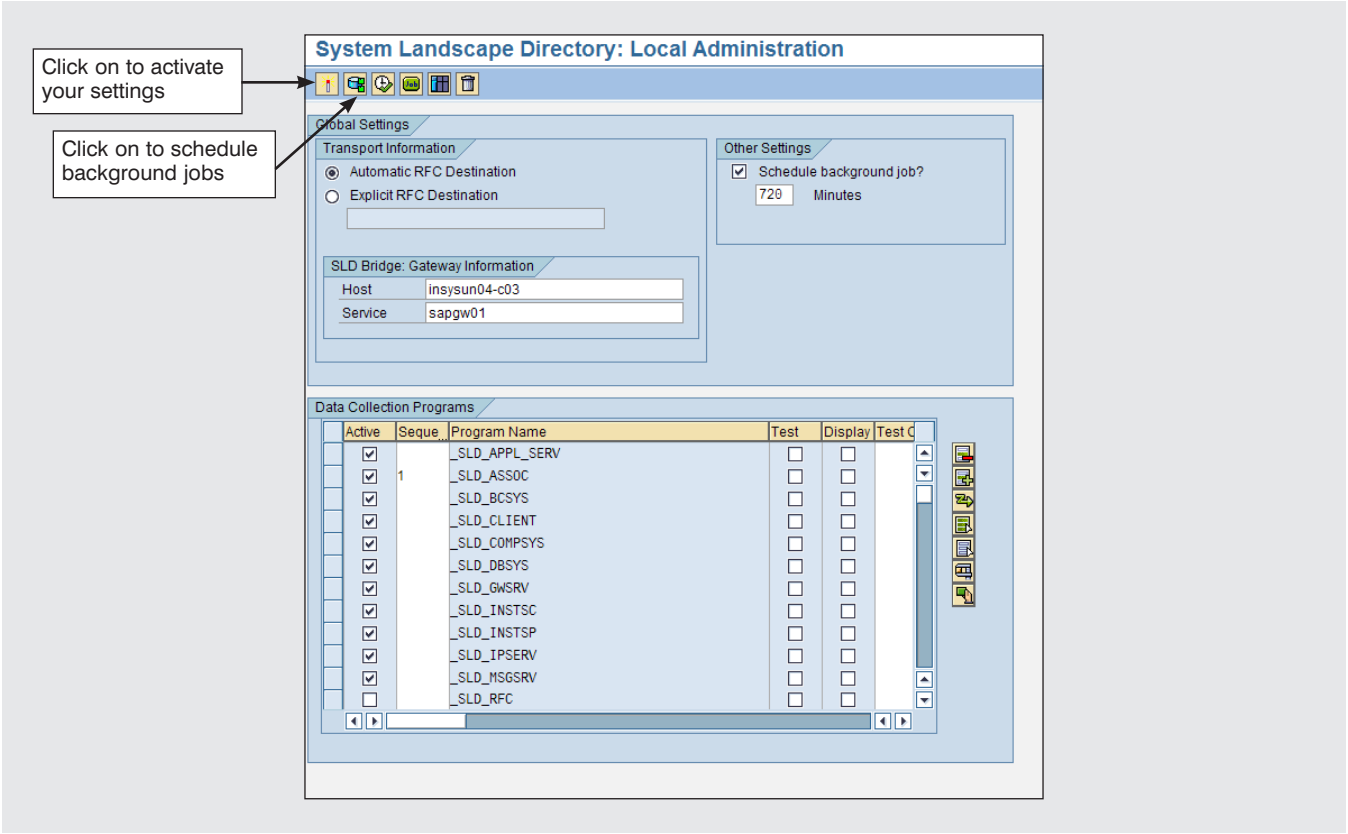

Figure 8 Setting up a data supplier

custom button instead of the typical button after starting wizard execution.

- **Multiple instancing:** If you need to run several instances on one host, you need a key from product management. You also need to sign a form in which you agree: running more than one instance on a host is at your own risk, and you need to consider certain side effects, such as what happens when configuring multiple database instances. For more information, see SAP Note 919135 – AC: Multiple Services on one Computing Node.

- **Security settings:** You can choose to configure the host agent credentials individually per host agent or to use one default account for all host agent connections. Then, you enter this default account under Default Credentials for the Host Agent.

- **User exits:** You can run custom code before stopping, after stopping, before starting, and after starting an instance with the ACC. If you would like to use user exits, check the checkbox.

- **Automatic mount points:** This is only interesting for large SAP hosting environments. If the creation of a necessary mount point has been missed on a host where the ACC starts a service, then the ACC can automatically create and delete the mount points on the host if you configure it to do so.

- **Timer settings:** You can increase or decrease the refresh times.

Whenever you want to change your parameters for adaptive computing, such as by modifying the refresh times for the ACC user interface (UI) or enabling or disabling multiple instancing, you need to re-execute the wizard. After the wizard re-executes successfully, you can use the undocumented config servlet (which you can find at http://<*acchost*>:5*nn*00/ACCConfig/BaseConfigurationWelcomeServlet) to modify the ACC parameters' values.

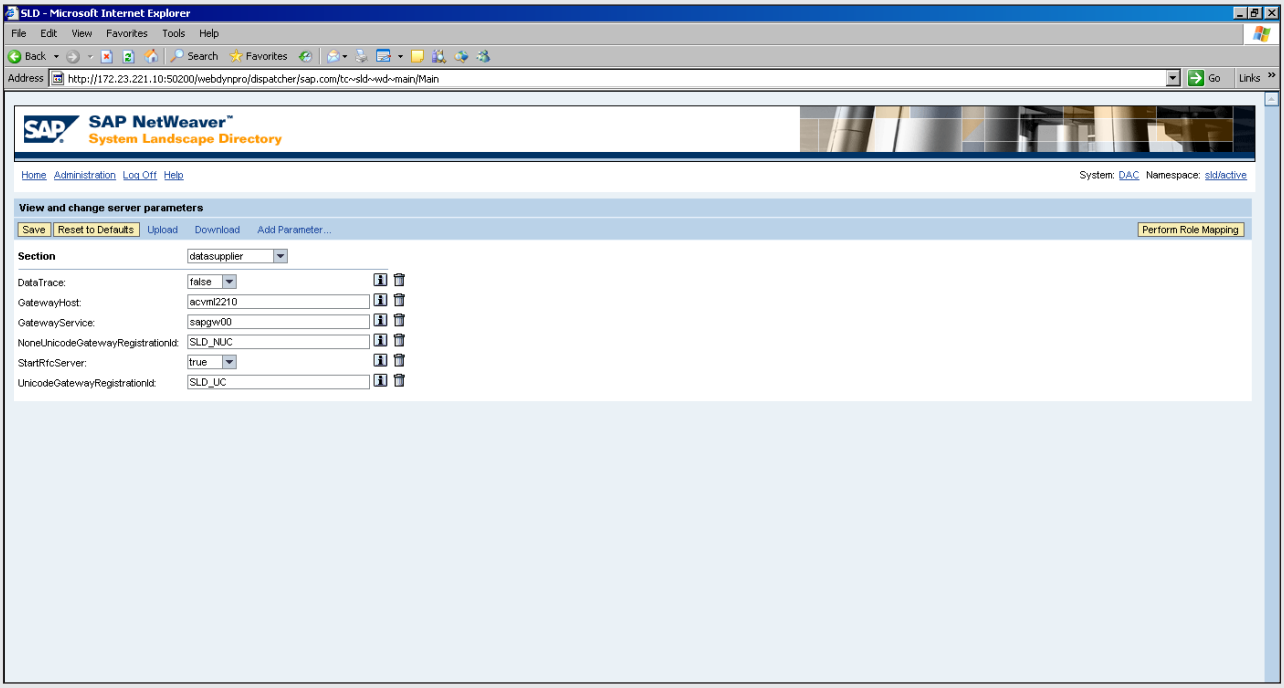

Figure 9 Choosing a gateway for ABAP-to-SLD communications

Report instances to the SLD

The data suppliers in your SAP systems, such as RZ70, collect and send data to the SLD. For example, RZ70 schedules jobs to collect SLD-relevant data and send it to the SAP gateway via RFC. The RFC then forwards this data to the SLD where the gateway is registered. After the ACC instance has been configured, the data suppliers update the SLD with reliable and up-to-date data at system startup or periodically by batch job.

You configure Java-based instances in the SLD data supplier section in J2EE Engine Visual Administrator. The Visual Administrator is a standard SAP tool for configuring certain runtime parameters of your J2EE engine. This tool is part of all J2EE instances in SAP Web Application Server (AS) 6.20 and 6.40 and SAP NetWeaver AS 7.0; however, in SAP NetWeaver AS 7.1 you use the NWA to configure instance-specific parameters. For ABAP-based instances, you use transaction RZ70 to set up a data supplier. First, you enter the host name and the port name of your installed SAP gateway. After you apply those values (see **Figure 8**), RZ70 schedules background jobs to collect instance-specific data and send it via RFC (this uses the TCP/IP connection `SLD_NUC`, `SLD_UC`, which RZ70 also configures depending on the entries you made) to the gateway. The gateway then forwards this data to the SLD. While Java instances can directly report to the SLD via HTTP, ABAP-based instances need the gateway.

Once you choose a gateway for ABAP-to-SLD communications (standard settings are shown in **Figure 9**), you enter it in the Administration section of your local SLD running on the ACC instance. Navigate to http://<*acchost*>:5*nn*00/sld, and then follow the menu path, Administration → Settings → Section data supplier. Then, enter the host (e.g., acvml2210); service (e.g., sapgq00), which is the name of one of the ports; and the registration information (e.g., SLD_NUC) for the gateway. An SAP

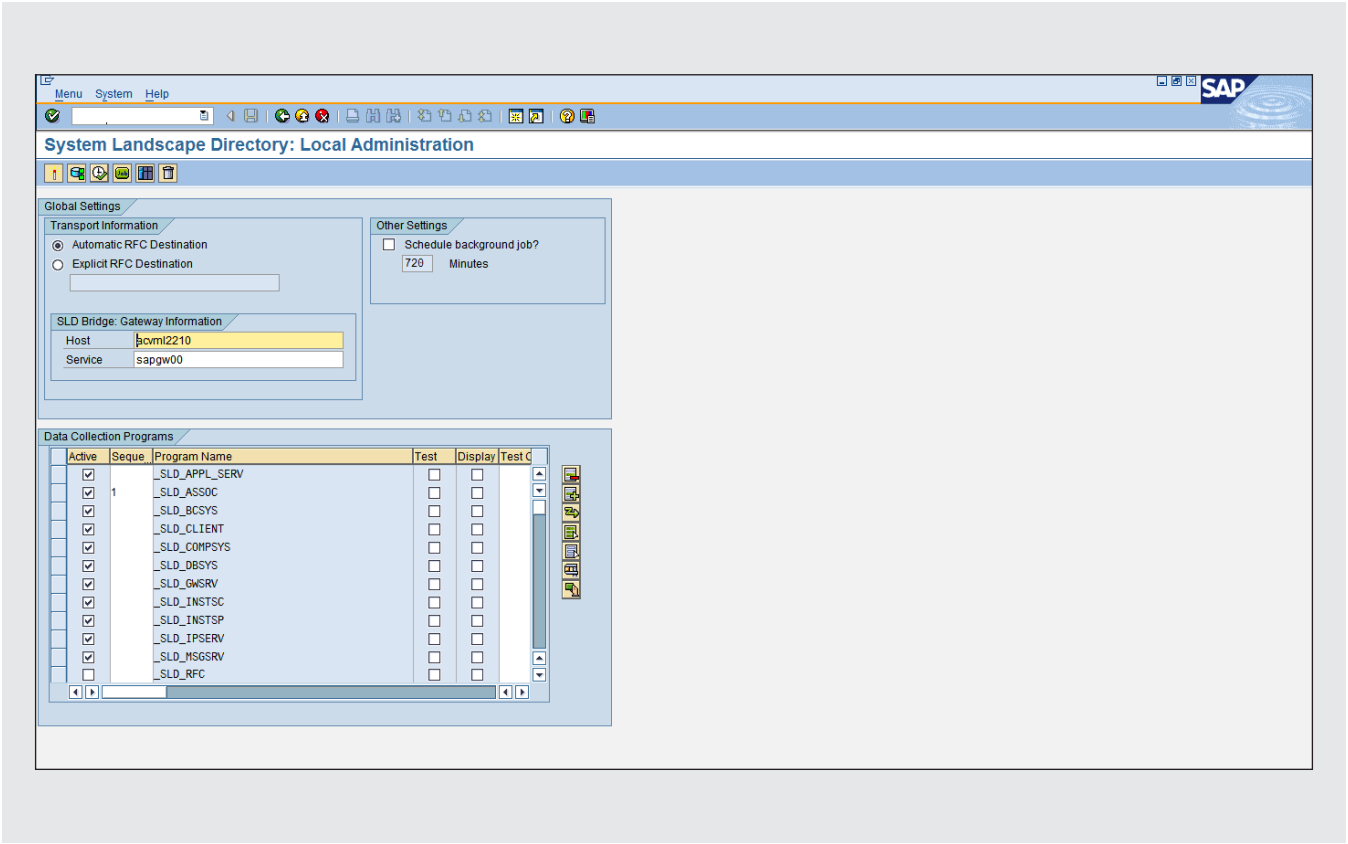

Figure 10 Registering an ACC-managed SAP instance in the SLD

gateway is always installed with a non-Unicode kernel. It will use the corresponding TCP/IP data provider depending on whether your application system is a Unicode or a non-Unicode system.

After you register the gateway, use transaction RZ70 to register your ACC-managed SAP instance in the SLD, as shown in **Figure 10**. Press F5 to schedule the data collector jobs, which collect the data to be sent to the SLD (e.g., host name of the instance, software component versions installed in the system, and kernel version). After the data collection process, look in your SLD to see whether the instance registered successfully, as shown in **Figure 11**.

You need to register all the application services in an adaptive landscape in your SLD. Make sure that the registered services run on those systems that have virtual host names.

Add a managed system to the NWA

Although the configuration guide doesn't mention it, you need to register all managed instances within the NWA to make them visible to the ACC; it's not enough to enter the SAP instance in the SLD. Click on the NWA Add Managed Systems task and press the Start or Re-execute button depending on whether you have already added a system, as in **Figure 12**.

You should execute this task after your instance reports to the SLD because then you can automatically fetch all the information that the wizard needs. When you execute this wizard, it automatically registers the instance in the NWA. The URL is neither confidential nor proprietary. First, the system must be registered in the SLD; then, you must execute this wizard to add it to the NWA.

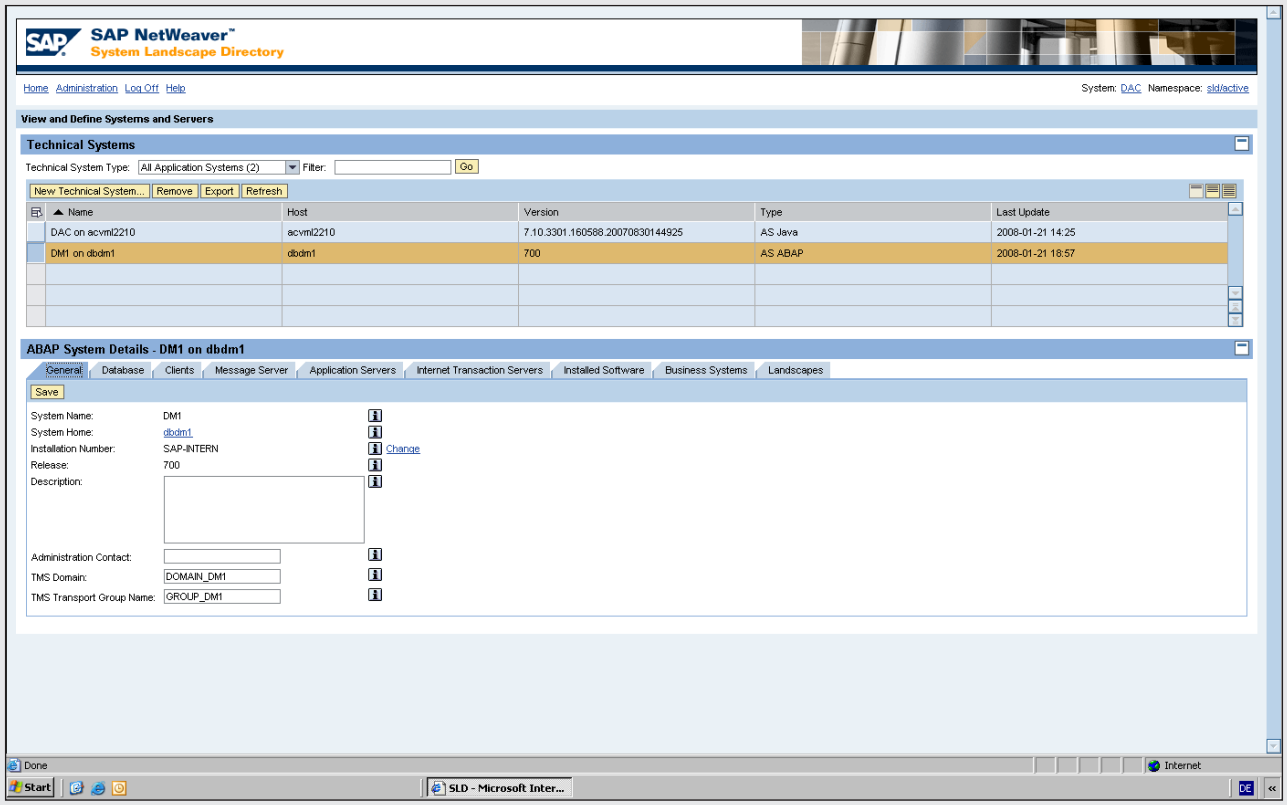

Figure 11 The SLD after data collection

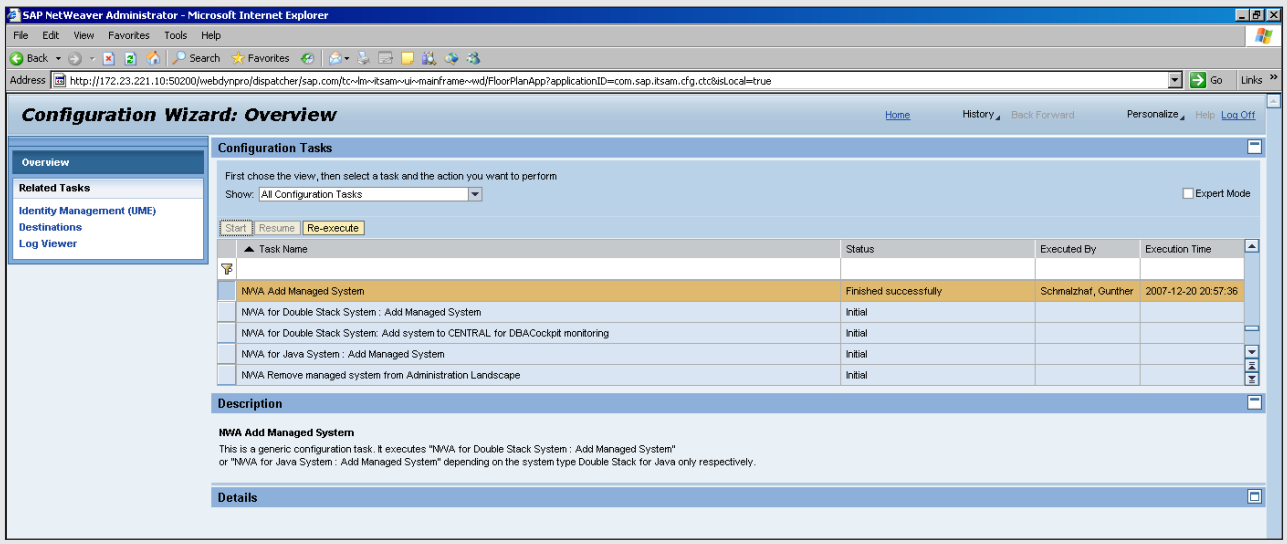

Figure 12 Adding a managed system to the NWA

Install the SAP host agents

The SAP start service provides monitoring and administration functionality and is one of the SAP NetWeaver management agents that you need to install on your computing nodes. It registers them to the SLD (see SAP Note 1031096 – Installing Package SAPHOSTAGENT for details). You can download the installation package from the SAP Service Marketplace's software distribution center at http://service.sap.com/swdc by navigating the following menu path: Download → Support Packages and Patches → Entry by Application Group → Additional Components → SAP Kernel → SAP KERNEL 64BIT/32BIT → SAP KERNEL 7.1 → <*platform*> → #database independent → SAPHOSTAGENT<*patch level*>.SAR.

The two types of SAP start services are host agents and instance agents:

- The *host agent* is part of the 7.1 kernel, and you must install it on every host that the ACC will manage. You can download the host agent from SAP Service Marketplace in the kernel update section for 7.1 (saphostagent.sar). The host agent offers a Web service interface that the ACC uses to connect to the host. Through the Web service, the ACC can then mount and unmount the storage resources on the host that belong to the instance it would like to start or stop on this node, as well as activating or deactivating the virtual IP addresses via sapacosprep (an executable that is also part of the host agent).
- The *instance agent* is already part of the kernel for every instance equal to or greater than SAP NetWeaver 7.0. For earlier versions, you can install the instance agent later (described in the ACC collective node). The instance agent is used to monitor the status of the instance (running, stopped, starting, stopping) and for starting and stopping an instance. It has a Web service interface, too. The new SAP Java management console (in a new layout like Microsoft Management Console's) also uses the Web service interface. Directly access the new Java MMC under http://<*hostname*>:5*nn*13 and use it for starting and stopping your instance or viewing instance log files or work process status. The MMC you've known for the Windows environment for years can now also be used for UNIX systems.

> ***Note!***
>
> For information on these services, see SAP Note 936273 – sapstartsrv for all platforms and SAP Note 1031096 – Installing Package SAPHOSTAGENT.

Saphostexec is part of the host agent and a UNIX daemon that executes an OS command, such as activating (or deactivating) virtual IP addresses, collecting OS statistics through saposcol, and authenticating users. Saphostexec runs under the context of a privileged OS user and controls all the functions that require a user of this type (e.g., acosprep, which activates and deactivates virtual IP addresses on its host).

Saphostctrl is a UNIX daemon that implements Web service runtime and interfaces. The host agents and the instance agents communicate with each other through UNIX domain sockets. After you download saphostagent.sar from the SAP Service Marketplace, create the user sapadm with root privileges and install the host agents by issuing the command `saphostexec -install` (see **Figure 13**) It will install them under `/usr/sap/hostctrl`.

To install the SAP host agent for the first time, you need to create a user sapadm who has group sapsys. You will later use the sapadm user to connect to the host agent (you must find some way to authenticate the ACC to the host agent because it provides critical functionality). The user sapadm doesn't need a logon shell or any profiles and can


```
- create user sapadm:sapsys
- mkdir -p /usr/sap/tmp
- Download saphostagent.sar from Service Marketplace and decompress in temp. dir.
- ./saphostexec –install
Install service
Check Installation Path ...ok
Copy Executable ...Thu Nov 15 02:21:26 2007
ok
Generate Profile ...ok
Install SYS-V service ...-> Start /usr/sap/hostctrl/exe/saphostexec
        pf=/usr/sap/hostctrl/exe/host_profile <-
Thu Nov 15 02:21:26 2007
start hostcontrol using profile /usr/sap/hostctrl/exe/host_profile
ok
```

Figure 13 Installing the SAP host agent for the first time on a computing node that the ACC manages

be a non-logon user (you may set logon/shell = false). Then, you need to extract saphostagent.sar to a temporary directory. From this directory, you execute `./saphostexec –install`, which then automatically installs the host agent under `/usr/sap/hostctrl` (a local directory).

After successfully installing the SAP host agents, you need to create an SLD configuration file. Within it,the host name and the SLD connection user (i.e., a user on the SLD J2EE Engine who has permission to write data to the SLD; the role `SLD_DATA_PROVIDER` allows the user to write data only to the SLD) are stored so that you can register the host within the SLD, as shown in **Figure 14** on the next page. You then have to manually register your instance within the SLD. You can use the following code snippet to do so.

```
cd /usr/sap/hostctrl/exe
./saposcol -b | ./sldreg -connectfile
          slddest.cfg
```

The documentation doesn't mention that you have to restart the host agents. After the first registration, you need to execute the commands shown in **Figure 15** (on page 45) to do it manually.

The host should now be visible in the SLD under Technical Systems (go to http://<*acchost*>:5*nn*00/sld and then navigate the menu path Landscape → Technical Systems. Choose Computer Systems from the Technical System Type drop-down list. Highlight the host name of the system for which you just installed the host agent (see **Figure 16** on page 45).

Now that you have successfully installed the ACC, the SLD is up and running the host agents of your computing nodes and the SAP instance report to the SLD. Let's look at how you add application services and computing resources so the ACC can manage them. Also, be sure that you have configured your instances for Web service authentication according to SAP Notes 927637 and 958253. You don't need to configure these instances if you already use Pluggable Authentication Modules (PAMs), as in Linux platforms. If not, those two notes describe the configuration, as do the ACC guides. Depending on your platform, you may need to edit `/etc/pam.conf` to enable PAM for service sapstartsrv. Under AIX, for example, I didn't have to edit `pam.conf` at all.

```
acvml2212:/usr/sap/hostctrl/exe # .sldreg -configure slddest.cfg

Thu Nov 15 02:37:05 2007 SLD Registration Program [2.0]  (Sep 30 2007)
     [Non-Unicode]

Use configuration file 'slddest.cfg' for output!
Please specify the HTTP connect information: [User, Password, Host, Port,
     HTTP/HTTPS]
User Name       : SLD_CI_DM1
Password        : ********
Verify Password : ********
Server Host     : acvml2210
Port [80,443,1024..655351]: 50200
Use HTTPS? (y/n)  : n

====================
Entered Information
====================
User:   SLD_CI_DM1
Passwd: (secret)
Host:   acvml2210
HTTPS:  n
====================

Write this information to secure file 'slddest.cfg'? [y=yes/n=no/q=quit] y

Thu Nov 15 02:38:42 2007 Writing data to encrypted file 'slddest.cfg'
Thu Nov 15 02:38:42 2007 Data written to secured data file 'slddest.cfg'
Thu Nov 15 02:38:42 2007 Using destination file 'slddest.cfg'.
Thu Nov 15 02:38:42 2007 No key file 'slddest.cfg.key' used.
Thu Nov 15 02:38:42 2007 Use encrypted destination file 'slddest.cfg' as
     data source
Thu Nov 15 02:38:42 2007 Configuration data successfully stored in file:
     slddest.cfg
acvml2212:/usr/sap/hostctrl/exe
```

Figure 14 Storing the instance to be registered in the SLD

Add instance and servers to the ACC

You can reach the start page of the ACC SP 4 (as shown in **Figure 17** on page 47) only with a URL (you will find the correct URL in the ACC 7.1 Collective Note 1008828 – ACC 7.1 PI/Adaptive Computing Controller Collective Note). In SP 5, SAP will have a quick link /acc.


```
acvml2212:/usr/sap/hostctrl/exe # ./saphostexec -stop
Thu Nov 15 02:44:23 2007

saphostexec is already running (pid=22727).
Stopping...Thu Nov 15 02:44:25 2007
Stopped

acvml2212:/usr/sap/hostctrl/exe# ./saphostexec pf=/usr/sap/hostctrl/exe/host_profile
Thu Nov 15 02:44:57 2007

start hostcontrol using profile /usr/sap/hostctrl/exe/host_profile
```

Figure 15 Restarting the host agents

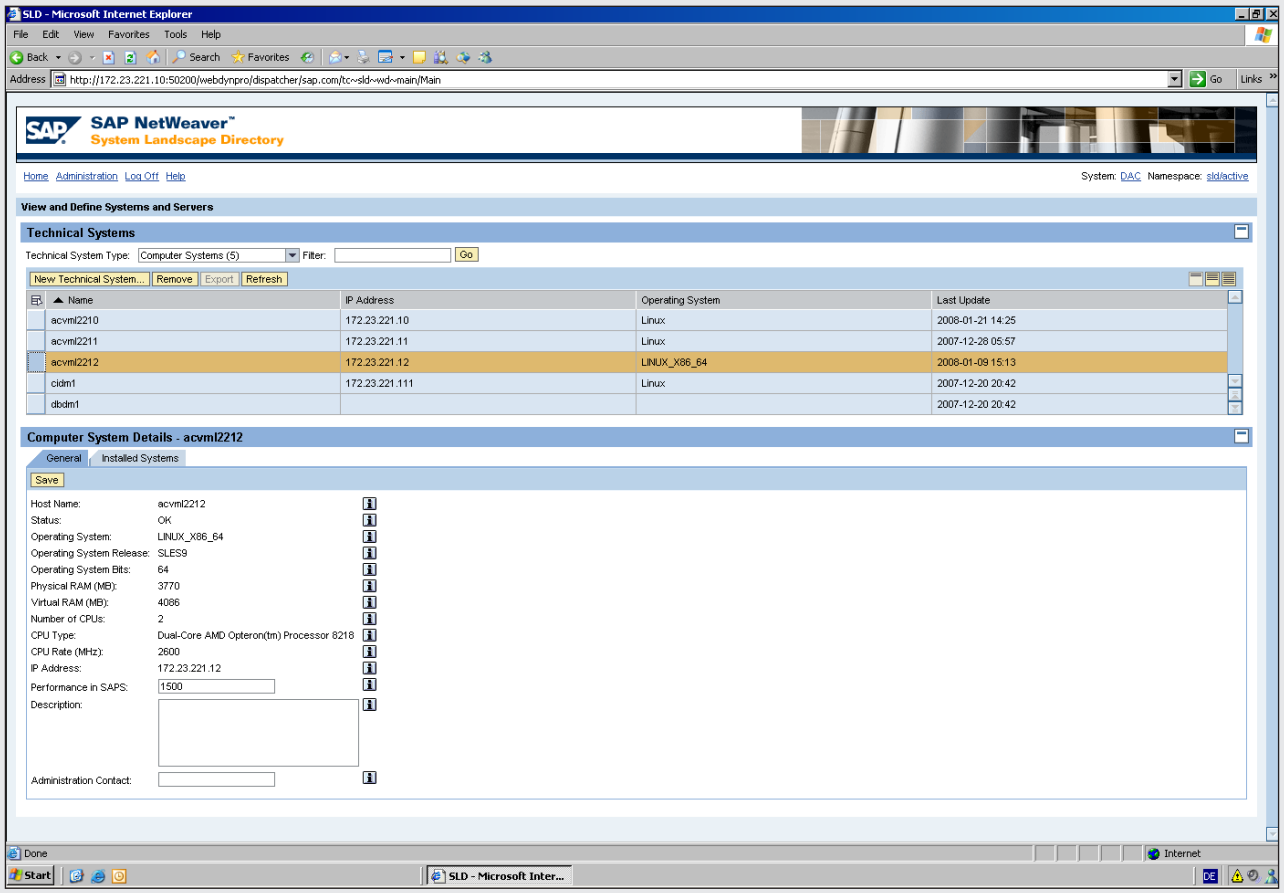

Figure 16 Note the host name of the system for which you just installed a host agent

Note!

PAM and Web service authentication correlate to each other. To make SAP host agent authentication possible, you need to configure PAM for your platform. For Linux, the following PAM configuration (extracted from SAP Note 1008828) will work.

On Linux you have to create a PAM authentication policy for the service sapstartsrv. The following example is a good starting point that should work for all Linux distributions. Create the file `/etc/pam.d/sapstartsrv` and add the following lines:

```
auth      sufficient  pam_unix2.so
auth      required    pam_unix_auth.so
account   sufficient  pam_unix2.s
account   required    pam_unix_acct.so
```

If you're running SUSE Linux Enterprise Server (SLES) 10, make sure you use "md5" as the default encryption method. See Note 958253 for more details.

Important!

Quick link /acc will be delivered with the next support package, SAP NetWeaver PI 7.1 Java SP 5, but did not make it into production for the current service pack, SP 4.

Configure the computing pools and networks

Instances and nodes go in an adaptive computing pool. You can create several pools for your environment, but instances that belong to one pool can only be moved to computing nodes assigned to the same pool. Also, an instance or node cannot belong to more than one pool at a time.

You can create a new pool by clicking on Pool & Network Configuration. After that, the configuration panel opens, allowing you to create new pools and networks. Then, click on the Pools tab (see **Figure 18**).

By clicking on the Networks tab (see **Figure 19** on page 48), you can configure one or more ACC networks. You must assign the network mask and the broadcast address of your server network here because the ACC needs this information to determine how to correctly activate the virtual host names on a server's NI. Later, you will assign your servers to an ACC network and also assign the instances to a network (or more than one if you have separated front-end and back-end networks). You can only relocate services on servers that belong to the same network as the services do.

Those networks will later be assigned to the services you configure for adaptive operation. You can give the network the name of your choice. Then, enter the network mask, which is dependent on your server network topology. You can easily determine the network mask by asking your network administrator or by looking at the interface configuration of your servers (enter `ifconfig` under UNIX/Linux). This is the subnet mask that determines to which subnetwork your IP address belongs. Then, you enter

Figure 17 Configuring the various elements of adaptive computing

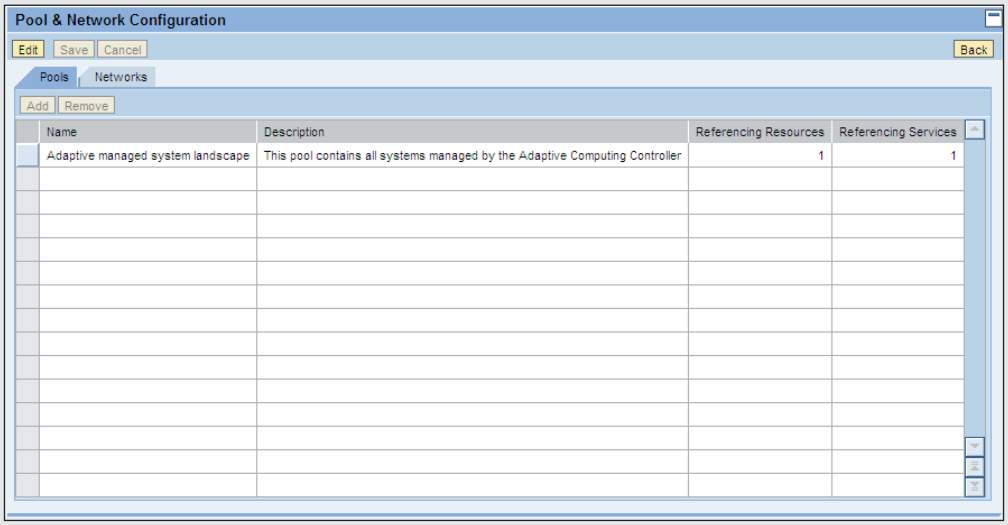

Figure 18 Creating a new adaptive computing pool

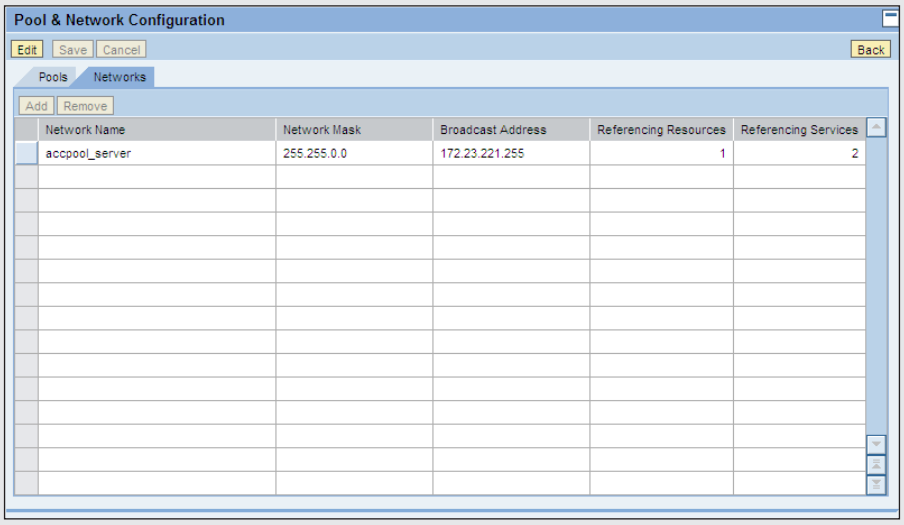

Figure 19 Configuring your ACC networks

the broadcast address of your network (depending on your network's topology).

Add computing resources to the adaptive computing pool

You can add computing resources by clicking on Resource Configuration (**Figure 17**). In this step, you add the servers to your adaptive computing pool, as shown in **Figure 20**.

The server with host name acvml2211 has already been configured, so now you want to configure the server with host name acvml2212. Highlight the server in the list, and click on the Edit button. This brings you to the configuration process.

In Step 1, as shown in **Figure 21**, you define the server to be managed in your adaptive computing environment by selecting the AC Managed checkbox. The server will then be visible in the ACC and you can see what SAP services are currently running on the server. You can also add servers here that run non-virtualized instances and that aren't supposed to be part of the adaptive computing landscape. (You can monitor the CPU utilization and memory consumption of those servers with the ACC, so it will give you a general overview of which systems you have and what the current utilization of those systems is.)

In Step 2, you specify the Resource Properties (see **Figure 22** on page 50). First, you identify the kind of adaptive enablement you want.

- **AC Enabled:** Select if you want to be able to start and stop services running on the server. If this is the only option you select, you cannot select this server as a relocation target in the ACC later. It will only allow you to start and stop services installed on the server and monitor its resource consumption (e.g., memory, CPU).
- **AC Operational:** Select to set the server for full integration in the ACC and to be able to relocate services to it.

In the Capabilities section, the number of available SAPS (a performance measure using the QuickSizer tool on SAP Service Marketplace) sets up an approximate number of SAPS for the server. If you try to relocate a service that needs more SAPS than the server can provide, you will see a warning displayed, and the service's auto-server-select functionality won't select that server.

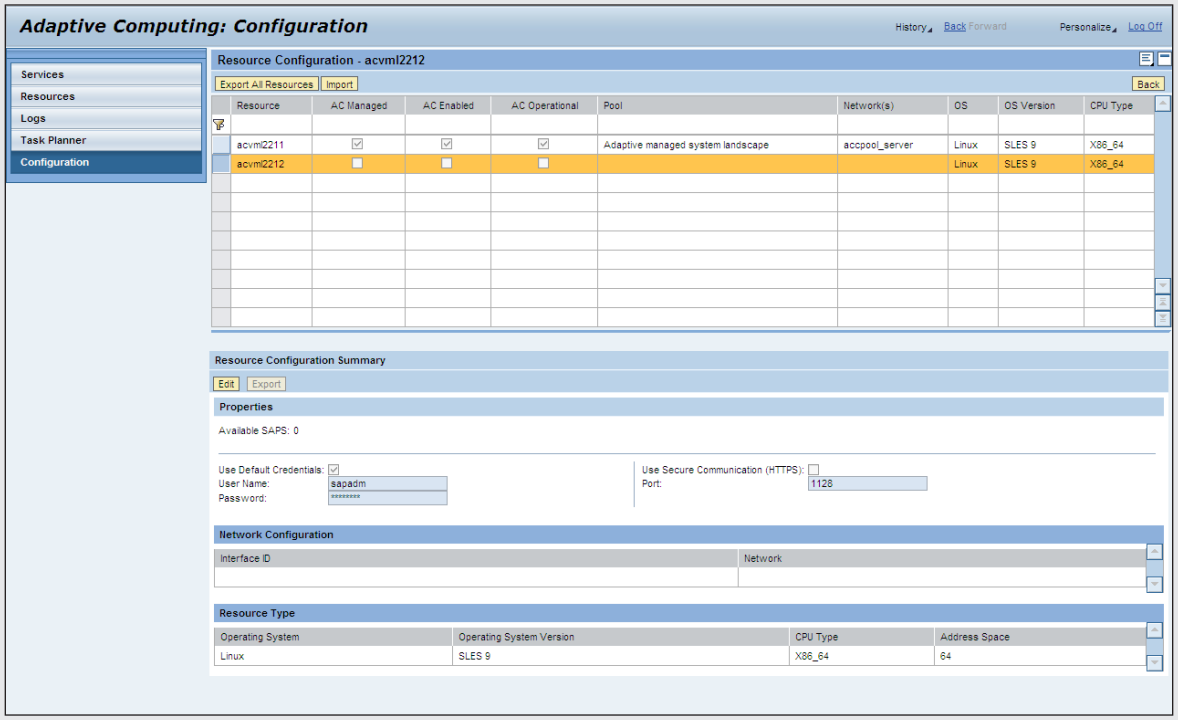

Figure 20 Adding computing resources to your adaptive computing pool

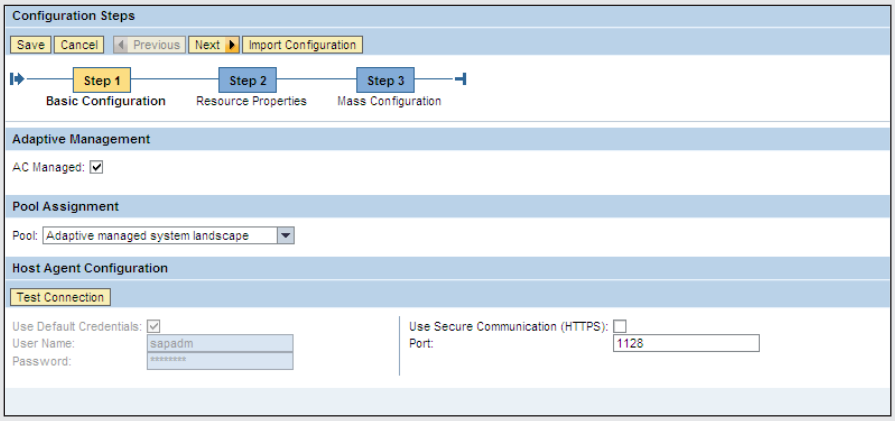

Figure 21 Defining the server that the ACC will manage

In the Network Assignment section, you assign the network that you previously specified (i.e., `accpool_server` to your server's corresponding NI ID (`eth0`). Later, when starting a virtualized

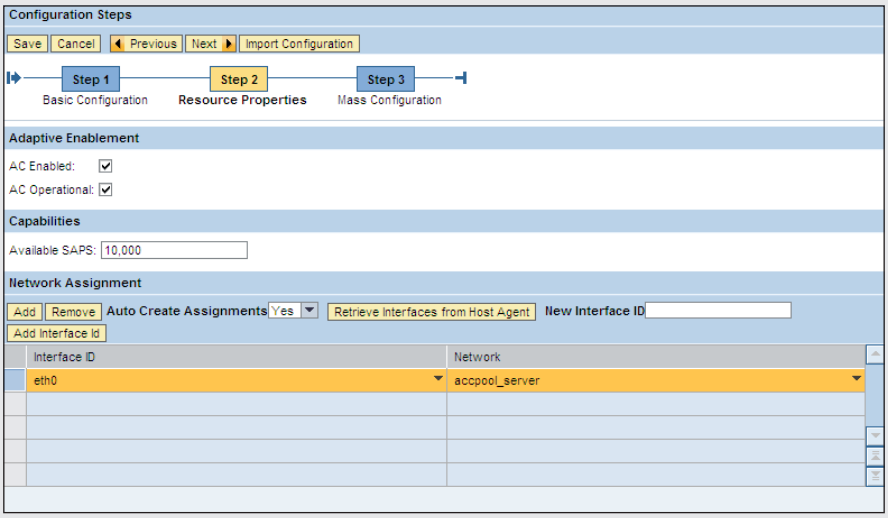

Figure 22 Allocating SAP service resources

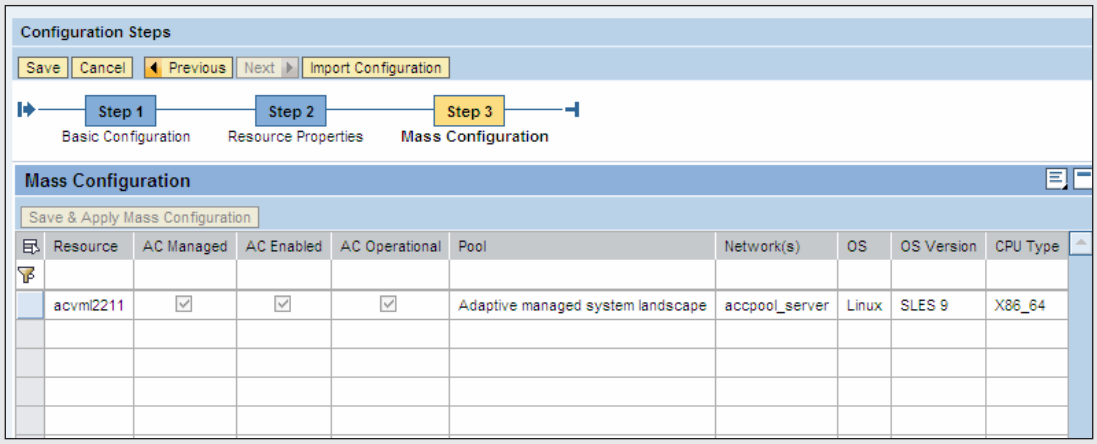

Figure 23 Applying configuration settings to other servers via mass configuration

application service here, it will activate the virtual IP addresses on this NI.

In Step 3 you can apply all the configurations you just made to any other server in your Access Control List (ACL), referred to as "Mass Configuration," as shown in **Figure 23**. For example, let's say you have 80 blade servers in your system landscape that you want to configure for the ACC. Instead of going through the configuration for every one of those 80 servers, you can start with the configuration of the first server. After you finish that one, the ACC will ask you if you would also like to apply the configuration you just made for this server to any other servers. You can then select the servers that need the same configuration values as your current server from the Mass Configuration lists.

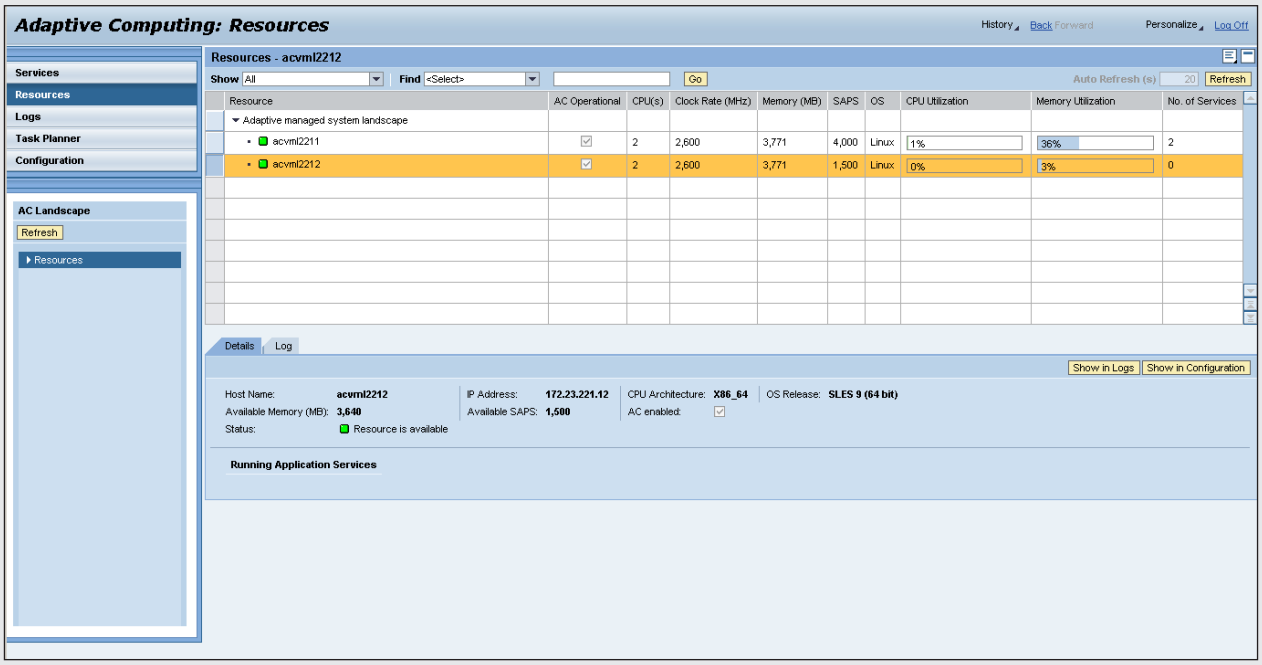

Figure 24 Displaying the newly configured server acvml2212 in the Resources overview

After a successful completion, you can see the server you just configured in the Resources overview (see **Figure 24**). Now, it's ready to run services.

Configuring application services

Next, you need to configure the application services. An application service can be an SAP CI, DB instance, DI, or a central service Java instance. To configure application services that report to the SLD and have been added as managed systems in the NWA, click on Service Configuration (**Figure 17**). Here, you configure an SAP NetWeaver AS ABAP system with system ID (SID) DM1, the SID of the CI. The DB instance has the same SID, of course. The SID consists of three characters, and the system was just installed on DM1. The DB instance was installed on virtual host name "dbdm1" and the CI on virtual host name cidm1 (as shown in **Figure 25** on the next page).

By clicking on the corresponding service, an Edit button will appear. Click on it to switch into edit mode, and you'll be able to configure the service (in this case, the CI or the DB instance).

In Step 1, you want to enable the AC Managed option (as shown in **Figure 26** on the next page), which allows you to configure the service for the ACC to manage. It also enables you to configure a service that isn't running on a virtual host name to start and stop — but not relocate — with that service. Then, you need to assign an adaptive pool (e.g., Adaptive managed system landscape). You can also test the connection to the instance agents. Since the instance agents are already part of the 7.0 kernel, you don't need to configure them. You can test to see if the connection is working physically. You'll use the same port here that you used for the MMC, that is, port 5*nn*13, because it's the port of the Web service.

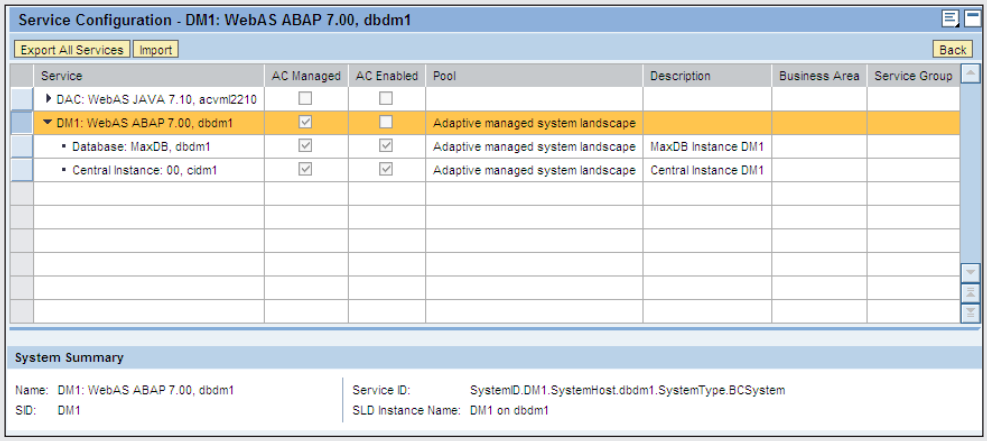

Figure 25 Specifying the application services to be configured for the adaptive computing landscape

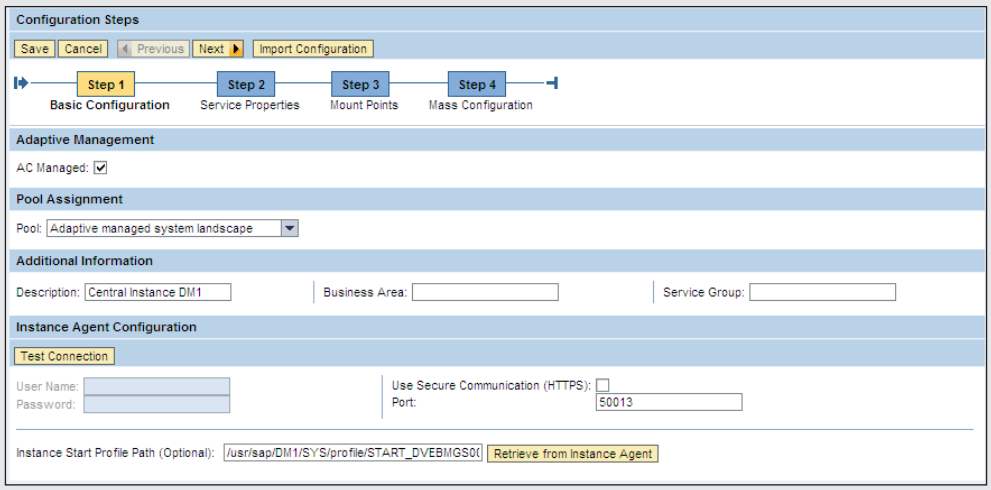

Figure 26 Specifying the application service for the CI DM1

In Step 2, you assign a network to the services virtual host name cidm1, as shown in **Figure 27**). If the service has AC Enabled turned on, the ACC can relocate it as long as it's running on a virtual host name. You can also enter the required number of SAPS (e.g., 500) and the amount of memory (e.g., 500MB) that the resource on which you start the service must provide. To ensure that you select the right server when you relocate a service, enter the appropriate values for your service here.

In Step 3, you configure the service's "mount points" (i.e., the local folder in which the remote folder's contents are stored). For a CI, you should configure the mount points shown in **Figure 28**

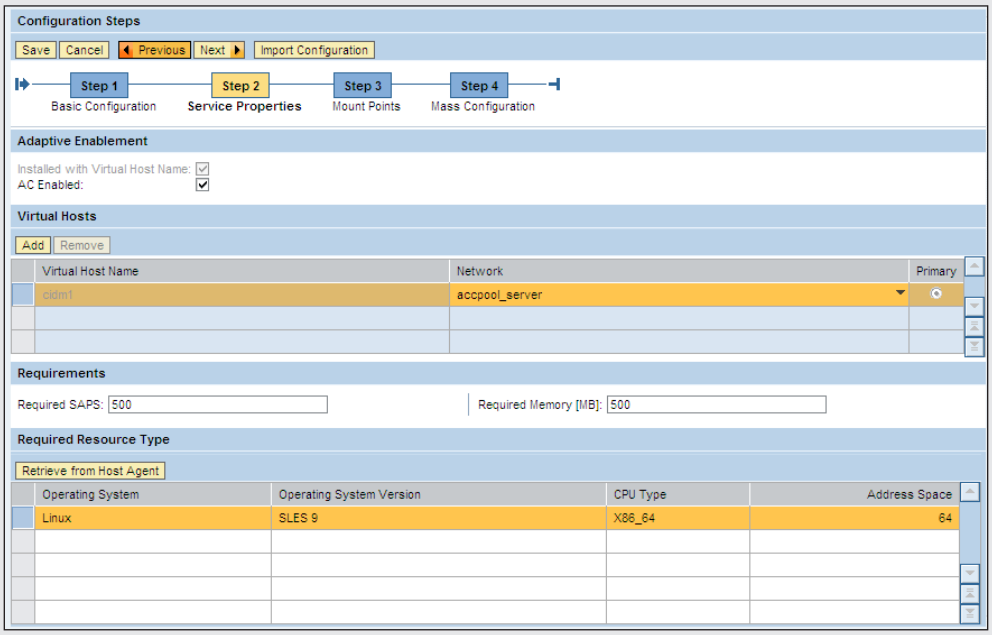

Figure 27 Assigning a network to the services virtual host name

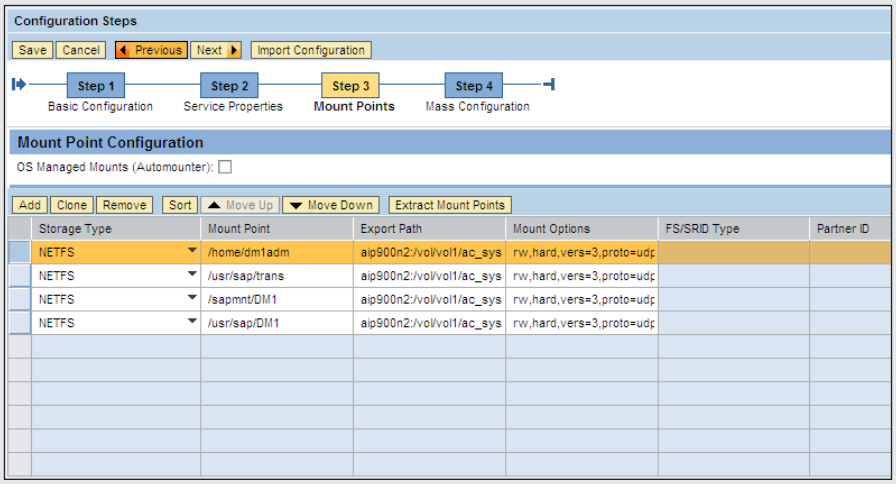

Figure 28 Specifying the mount points of the service for a CI

at a minimum. The example uses NFS-based storage and exports the NFS shares to the computing nodes. The ACC supports NFS storage technology out-of-the-box and needs no

Tip!

For NFS shares, the following mount point options are recommended but should also be based on your system landscape. Contact your storage hardware partner for a more detailed recommendation.

Option	Meaning
rw	Reads and writes allowed
nointr	Executable continues running on the client after a failover of the server occurs
hard	Client continues to attempt the NFS file operation indefinitely if operation fails
suid	Allows use of Set User ID (SUID) privilege on any files within the file system
vers=3	Version 3 of the NFS file system
proto=tcp	Protocol = TCP
lock	Automatically set if you do not explicitly set nolock as an option (This option guarantees proper handling of file locks on shared NFS files.)
rsize=32768	Size of NFS packages transferred when reading
wsize=32768	Size of NFS packages transferred when writing

other storage libraries. You can also extract the mount points directly from the instance or by copying the output of the `mount` command to the opening mount-point window by clicking on the Extract Mount Points button.

Tip!

For a MaxDB instance, you should also add mount point `/var/spool/sql` so you don't run into problems while relocating. In addition, make sure you synchronize file `/etc/opt/sdb` across all computing nodes.

In Step 4, you apply the configurations to additional application services through mass configuration. You also add the mount points of your database location here, as well as in the home directory of sdb (the administrative username of a MaxDB instance). The sdb username will be created during the database installation of a MaxDB; if you use a different database, then sdb will not exist as a user. This process is comparable to the one for configuring the CI except that the mount-point details differ.

After the final configuration, switch to the Services section of the ACC, which gives you an overview of all the configured application services (see **Figure 29**).

DM1 is now in the list of configured services and is running on server acvml2211. By selecting both of the dependent application services in the Mass Operations section, you can select the corresponding action in the Operation Type drop-down list, as shown in **Figure 30**, and execute the selected action (see **Figure 31** on page 56 for the available actions).

It's possible to operate your DB instance and CI on different servers, but further configurations of your system might become necessary. For example, you can't distribute instances if you haven't installed them as the distributed installation type during the SAP system setup. You can also run more than one instance on a computing node, but if you do, you will have to apply for a special key from the ACC's product management (they will check to see whether you fulfill the requirements for it).

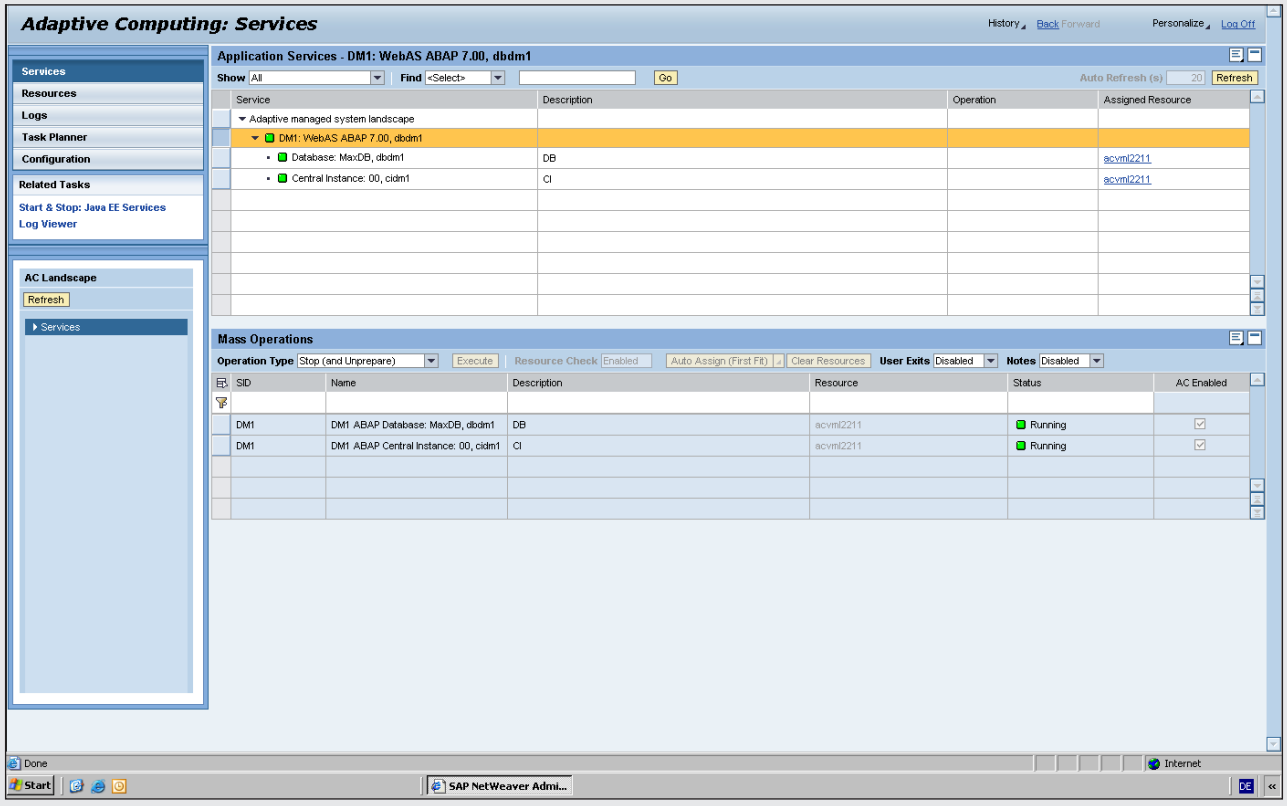

Figure 29 Summarizing all the configured application services

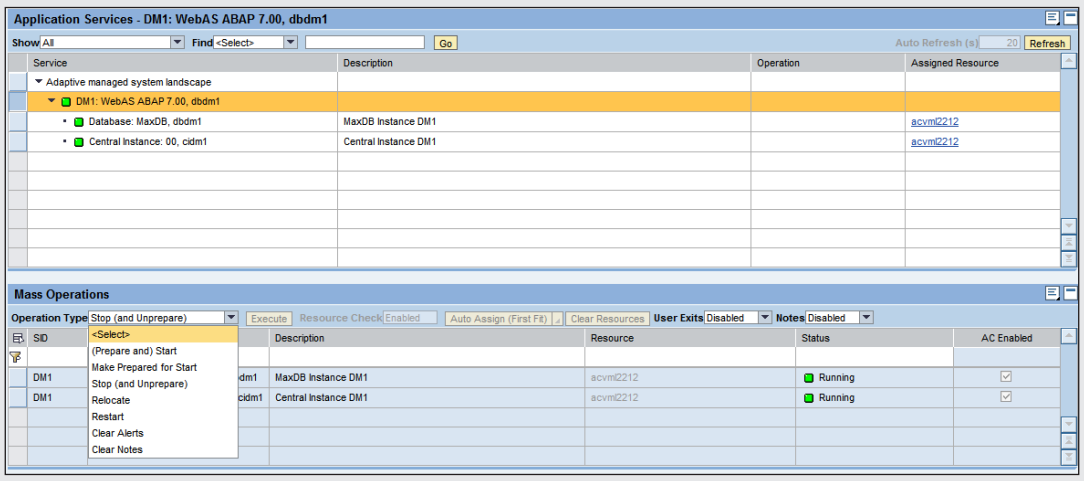

Figure 30 Operation Type drop-down menu (upper left of the Mass Operations dialog) shows all possible operations for this service

Operation type	Description
(Prepare and) Start	Activate virtual IPs and mount NFS shares to selected server; start instance here.
Make Prepared for Start	Only activate virtual IPs and mount NFS shares to selected server without starting instance.
Stop (and Unprepare)	Stop instance, unmount NFS shares, and deactivate virtual IPs.
Relocate	Move the application service to the selected server.
Restart	Stop and start service.
Clear Alerts	If something went wrong during the last operation, you can clear possible alerts here.
Clear Notes	If you set notes to an instance, clear them here. For example, you might need to shut down your system and set a note for ACC operators that says "System intentionally shut down, please do not restart." You can use this operation type to clear such notes.

Figure 31 Descriptions of the operation types available for application services

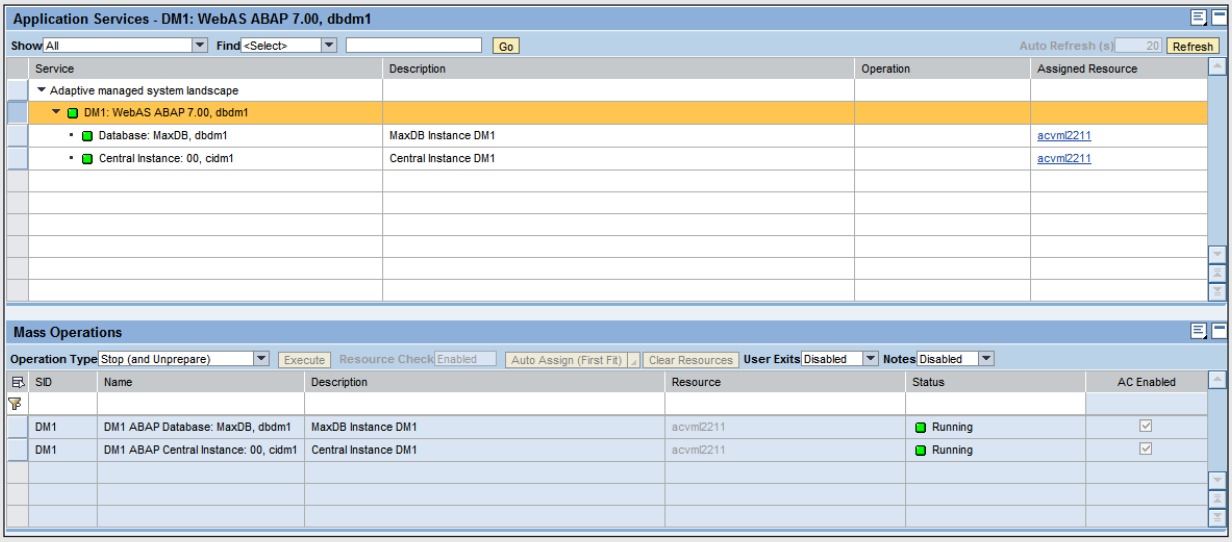

Figure 32 The application services overview showing the successfully relocated DM1

After the installation completes, you need to check to see whether the system is up and running. Depending on the version and platform you're running, you could simply try to connect to the system with SAP GUI/Web browser. If the connection succeeds, the system is up. Or you could use SAP-provided monitoring tools such as MMC, dpmon, and jcmon.

Now, let's look at an example that shows you what you can do (or what you have to do) when moving a system with the ACC.

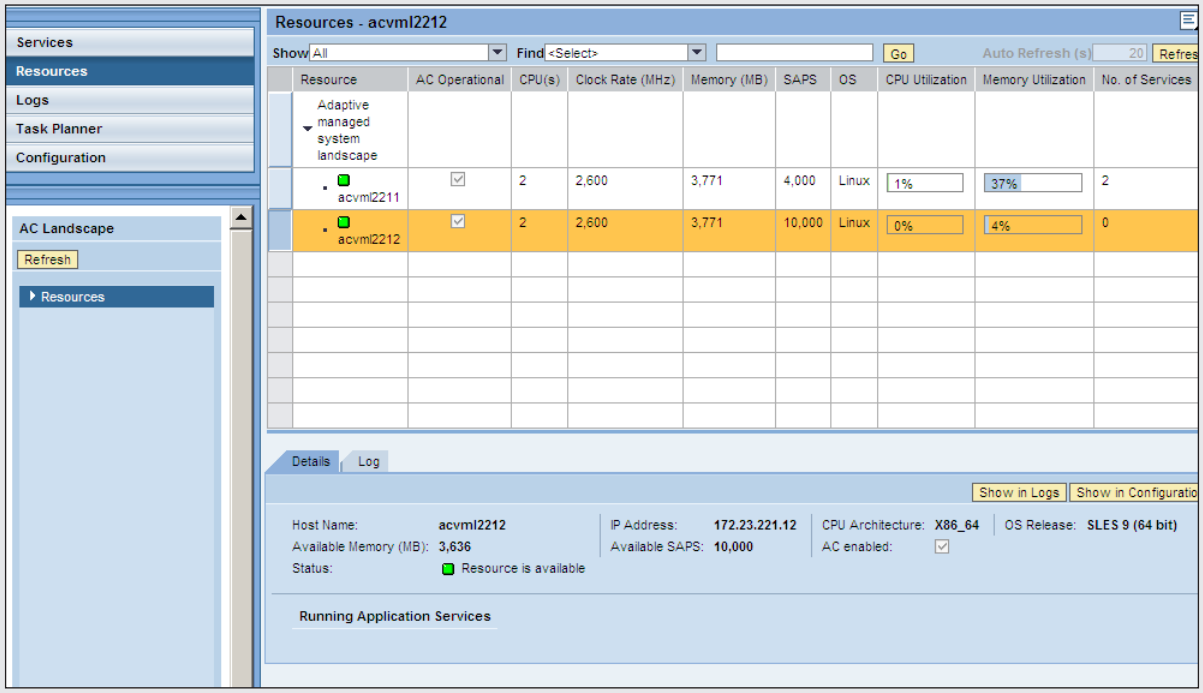

Figure 33 Server acvml2212 has no active services running, while acvml2211 has two

How to move a system with the ACC

The application service DM1, which consists of a database and a CI, currently runs on a server with virtual host name acvml2212. Let's move the service to a different server, perhaps because you need to upgrade the server's host OS,[1] you want to extend the server's memory, or you just want to run the service on a stronger hardware platform.

By clicking on the Operation Type drop-down arrow at the upper left of the Mass Operations dialog (**Figure 30**), you can see all possible operations for this service (**Figure 31**). Since you want to relocate to a different server, choose the Relocate option.

The system then asks you to which resource you want to relocate that application service. Since the server acvml2211 is a free server in the environment, relocate the service to the dedicated server (acvml2212). You can then schedule the service to stop. After the CI and the DB instance have been successfully stopped, they will be scheduled to start on target server acvml2211, as shown in **Figure 32**. After a couple of minutes, you can see the success of the operation in the application services overview. The ACC has successfully relocated service DM1 to the new target server.

Another check in the resource view (click on Resource Configuration shown in **Figure 17**) reveals that no more services are running on server acvml2212 and it is now free to start other services or run maintenance operations (see **Figure 33**).

[1] In a mature ACC environment, the OS of a server should always be deployed via Netboot or a central deployment function at server start. This will give you more flexibility when extending your landscape with new hardware since a predefined image will be used for all servers.

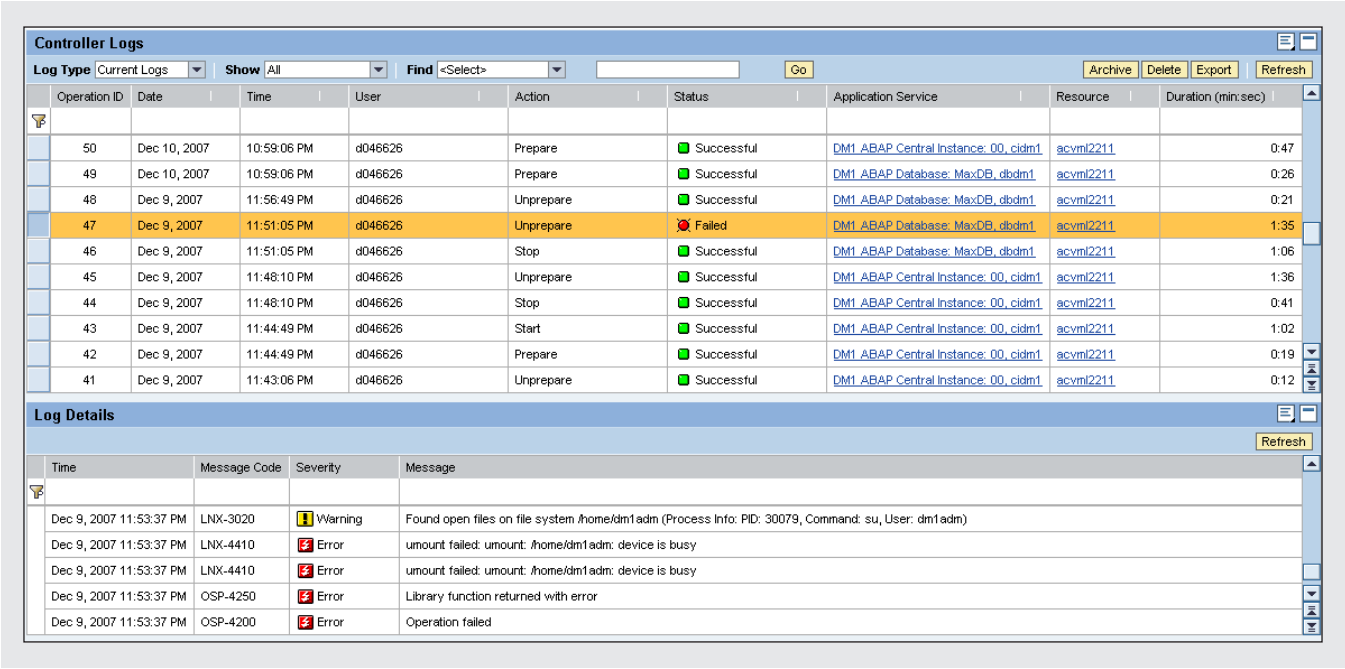

Figure 34 Listing the Controller Logs containing the actions that the ACC has performed

ACC security

Since access to the ACC means full control over its managed application landscape, you can start, stop, and relocate every system that the ACC manages. That means anyone who can access the admin user of your ACC can shut down any system attached to it because he or she won't be asked for the sidadm password when starting or stopping a system. To deal with this vulnerability, SAP has introduced the following new roles:

- **ACC_Observing:** The observer role bars these users from performing configuration or operation tasks. They see only the managed landscape and services, and they can't perform any actions, such as task planning or relocation of services.
- **ACC_Operation:** The operator role gives you the necessary rights for exception handling (e.g., resetting a service status) and operating the environment, but no rights to change configuration.
- **ACC_Configuration:** The configurator role covers configuration changes (e.g., adding or removing services or servers).
- **ACC_Archiving:** The archiver role lets you archive the ACC logs since they can grow rapidly and, therefore, consume a lot of valuable disk space in a large system landscape.

All actions that the ACC performs are written to the Controller Logs files (see **Figure 34**), which you can view by clicking on Logs in the left navigation panel.

The Log Details portion of the screen shows an error that can occur when someone is in a directory at the time the ACC tries to unmount it. The system then returns a "device is busy" error. In this example, I had a terminal window opened, was connected to machine acvml2211, was logged in as dm1adm, and was in the directory /home/dm1adm.

```
acvml2211:~ # su - dm1adm
Directory: /home/dm1adm
acvml2211:dm1adm 51> pwd
/home/dm1adm
acvml2211:dm1adm 52>
```


Since I was in the home directory, it couldn't be unmounted. To execute an unmount command, which the ACC triggers, I would need to leave the home directory or close the terminal connection. How do you react to this kind of error? First, try to determine who is currently in the directory and ask him or her to step out, or try a forced unmount depending on your OS. Then, you can reset the status of the application service via Clear Alerts (**Figure 31**). You can also try to manually unmount all NFS shares and manually deactivate the virtual IP addresses. Then, simply restart the service or bring it up on a new node by attaching the NFS shares and IPs and running the script `startsap all`.

Conclusion

This article shows you how adaptive computing can help you to improve the flexibility of your SAP solution landscape and how you can use the new ACC 7.1 as a central management tool for more flexibility, become independent of your hardware, run your instances on any free server, and move them whenever necessary. It also shows you how you can shut down any systems you only need from time to time and keep them in central storage where they no longer allocate servers.

It will take some time to fully introduce the ACC to your system landscape. Central user management and central OS deployment are some of the topics you must consider in the context of a global rollout of the ACC. Imagine a host environment such as the SAP Hosting organization, which manages hundreds of systems with ACC. The advantage of a central OS deployment is that you only have to patch one deployment image and the patch will become active on every server after a reboot. For the ACC, you can easily integrate new servers into your landscape without an OS installation.

After the hardware is up and running, the OS will be delivered via central OS deployment. Central user management is necessary to ensure that the users can logon to every node (such as sidadm, you never know to which servers you might relocate). You only need to manage one profile when you change the profile of a sidadm user, for example. If you have a local user environment, then changing the profile of a user would mean changing the user's profile on every server.

Perhaps the wisest course of action is to start small and grow carefully by virtualizing existing application servers (i.e., reinstalling them with virtual host names and putting them in central storage), and bit by bit shut down old application servers that use physical addresses. You can then continue to make CIs and DB instances virtual. In the end, you introduce multiple instancing by running several services in parallel on one host only. Select your instance numbers carefully since the port numbers depend on your instance number, something that could create problems when running different instances that have the same instance number on the same server.

Adaptive computing is part of the strong efforts SAP is continuously investigating to get the maximum value out of its technology and hold down total cost of ownership (TCO) in operating business solutions.

SAP has completely rewritten the new ACC 7.1. It's now based on Web services. In the future, virtual machine management will find its way into the ACC software. Without SSH communications, it will also be easy to manage Windows environments in the near future (SAP host agents for Windows and Oracle are currently under development).

The ACC gives you the opportunity to decouple your SAP applications from their underlying hardware. You can react to changing needs within minutes by starting additional application servers on free hardware or moving a system if the underlying hardware is broken or needs to be fixed. The ACC means enhanced availability for your SAP solutions. Improve server utilization by starting additional application instances on any free servers and shutting down systems not needed over the weekend. Then, use the freed resources to start additional application instances that support your other systems with batch work processes. It becomes easy to set up and integrate new hardware into your existing infrastructure. Once the hardware is in, move your system with the ACC to the new box. The flexibility to assign and make use of other hardware resources increases, too.

Adaptive computing makes the difference because every service is now available on each node of your network. High availability typically requires additional expensive hardware and increases complexity. With ACC, an increased availability is fundamentally included. Use computing power and storage as shared resources for growing, dynamic business needs. Satisfy your customers with higher service levels and reliability. And easily manage your system landscape since the ACC hides the complexity behind it.

The challenge for IT to generate benefits, and not costs, will increase in the future. The money you save can be invested in innovations instead of operations. In one customer's TCO examination, for example, savings from using the adaptive computing architecture are as much as 20%. An additional 5% to 10% can be made by using automated operations (e.g., use the Task Planner to automatically shut down training systems over the weekend, and use free resources to support financial and BI reporting systems). You can also expect an additional 5% to 10% savings through synergy optimizations. So, overall, you can save 30% to 40% with a strong ACC concept depending on the current state of your system landscape. It also increases your speed of reaction when business requirements change, which is measurable in monetary value — but also in enhanced service times and user satisfaction.

