



Migrating to the Solaris™ Operating System: Migration Strategies

Ken Pepple, Brian Down, David Levy

Sun BluePrints™ OnLine—September 2003



<http://www.sun.com/blueprints>

Sun Microsystems, Inc.
4150 Network Circle
Santa Clara, CA 95045 U.S.A.
650 960-1300

Part No. 817-4102-10
Edition: September 2003

Copyright 2003 Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, California 95045 U.S.A. All rights reserved.

This product or document is protected by copyright and distributed under licenses restricting its use, copying, distribution, and decompilation. No part of this product or document may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any. Third-party software, including font technology, is copyrighted and licensed from Sun suppliers.

Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California. UNIX is a registered trademark in the United States and other countries, exclusively licensed through X/Open Company, Ltd.

Sun, Sun Microsystems, the Sun logo, Sun BluePrints, and Solaris are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and other countries. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the US and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

The OPEN LOOK and Sun™ Graphical User Interface was developed by Sun Microsystems, Inc. for its users and licensees. Sun acknowledges the pioneering efforts of Xerox in researching and developing the concept of visual or graphical user interfaces for the computer industry. Sun holds a non-exclusive license from Xerox to the Xerox Graphical User Interface, which license also covers Sun's licensees who implement OPEN LOOK GUIs and otherwise comply with Sun's written license agreements.

RESTRICTED RIGHTS: Use, duplication, or disclosure by the U.S. Government is subject to restrictions of FAR 52.227-14(g)(2)(6/87) and FAR 52.227-19(6/87), or DFAR 252.227-7015(b)(6/95) and DFAR 227.7202-3(a).

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

Copyright 2003 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, Californie 95045 Etats-Unis. Tous droits réservés.

Ce produit ou document est protégé par un copyright et distribué avec des licences qui en restreignent l'utilisation, la copie, la distribution, et la décompilation. Aucune partie de ce produit ou document ne peut être reproduite sous aucune forme, par quelque moyen que ce soit, sans l'autorisation préalable et écrite de Sun et de ses bailleurs de licence, s'il y en a. Le logiciel détenu par des tiers, et qui comprend la technologie relative aux polices de caractères, est protégé par un copyright et licencié par des fournisseurs de Sun.

Des parties de ce produit pourront être dérivées des systèmes Berkeley BSD licenciés par l'Université de Californie. UNIX est une marque enregistrée aux Etats-Unis et dans d'autres pays et licenciée exclusivement par X/Open Company Ltd.

Sun, Sun Microsystems, le logo Sun, Sun BluePrints, et Solaris sont des marques de fabrique ou des marques déposées, ou marques de service, de Sun Microsystems, Inc. aux Etats-Unis et dans d'autres pays. Toutes les marques SPARC sont utilisées sous licence et sont des marques de fabrique ou des marques déposées de SPARC International, Inc. aux Etats-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc.

L'interface d'utilisation graphique OPEN LOOK et Sun™ a été développée par Sun Microsystems, Inc. pour ses utilisateurs et licenciés. Sun reconnaît les efforts de pionniers de Xerox pour la recherche et le développement du concept des interfaces d'utilisation visuelle ou graphique pour l'industrie de l'informatique. Sun détient une licence non exclusive de Xerox sur l'interface d'utilisation graphique Xerox, cette licence couvrant également les licenciés de Sun qui mettent en place l'interface d'utilisation graphique OPEN LOOK et qui en outre se conforment aux licences écrites de Sun.

CETTE PUBLICATION EST FOURNIE "EN L'ETAT" ET AUCUNE GARANTIE, EXPRESSE OU IMPLICITE, N'EST ACCORDEE, Y COMPRIS DES GARANTIES CONCERNANT LA VALEUR MARCHANDE, L'APTITUDE DE LA PUBLICATION A REpondre A UNE UTILISATION PARTICULIERE, OU LE FAIT QU'ELLE NE SOIT PAS CONTREFAISANTE DE PRODUIT DE TIERS. CE DENI DE GARANTIE NE S'APPLIQUERAIT PAS, DANS LA MESURE OU IL SERAIT TENU JURIDIQUEMENT NUL ET NON AVENU.



Please
Recycle



Adobe PostScript

Migrating to the Solaris™ Operating System: Migration Strategies

Editor's Note - This article is the complete third chapter of the Sun BluePrints™ book, *Migrating to the Solaris OS*, by Ken Pepple, Brian Down, and David Levy, which will be available from www.sun.com/books, amazon.com, and Barnes & Noble bookstores at the end of October, 2003.

This article defines the most important terms in migration and differentiates between these terms. In addition, it presents migration strategies, the benefits and risks of each strategy, and the appropriateness of each strategy for various situations.

This chapter contains the following sections:

- “Concepts” on page 1
- “Evaluating the Environment” on page 8
- “Strategies” on page 11
- “Choosing a Strategy and Developing Tactics” on page 18

Concepts

Within the context of this book, the term migration is defined as the transition of an environment's people, processes, or technologies from one implementation to another. While somewhat open ended, this definition allows us to discuss migration in a number of different contexts:

- The migration of common off-the-shelf (COTS) software from one platform to a larger or smaller similar platform

- The migration of data from one database to another, possibly similar data storage technology
- The migration of a custom-written application from one platform to a different platform and operating environment

These examples have different inputs, different execution strategies, and dissimilar functional outcomes, but they are all migrations in that the IT functionality has been moved from one platform to a different platform. It is this varied scope that affords the term “migration” its varied interpretations.

Consolidation

Migration is often confused with the act of *consolidation*. *Webster's College Dictionary* defines consolidation as the act of bringing together separate parts into a single or unified whole. As defined in the Sun BluePrints book *Consolidation in the Data Center*, by David Hornby and Ken Pepple, consolidation should be thought of as a way to reduce or minimize complexity. If you can reduce the number of devices you have to manage and the number of ways you manage them, your data center infrastructure will be simpler. This simplicity should contribute to efficiency, and consistency should contribute to lower management costs in the form of a reduced total cost of ownership (TCO).

While consolidations involve migrations, as applications and business functionality are moved to a single machine, migrations do not necessarily involve consolidation. The value proposition realized by migration relates to improved quality of service (QoS) and reduced TCO realized as a characteristic of the new platform, environment, and overall IT infrastructure.

Adoption

Frequently, the release of a new edition of an operating environment might require that the platform's operating system (OS) be replaced with a more up-to-date version, a practice referred to as *adoption*. Depending on the changes introduced in the new version, upgrading to a new version of the current OS might be as difficult as migrating the application to a completely different OS, even though the hardware platform remains the same.

Usually, applications that must be moved to a newer version of the same environment, have to be tested to determine whether they provide the same functionality in the new environment. This can be a time consuming, as well as expensive, process. If test suites designed to verify critical application functionality are not available, they will need to be developed, at considerable time and expense.

Application programming interfaces (APIs) are the touch points between applications and the operating environment. In addition to defining APIs, the Solaris environment also supports the concept of the application binary interface (ABI), the description of these APIs used by the application executable at a binary level. This definition enables you to compare the API usage of an application executable created in one version of the OS to the binary description of interfaces supported by a different release of the OS. Consequently, compatibility between different versions of the OS can be guaranteed without any examination of the source code used to create the application.

In most cases, binaries created under an earlier version of the OS require no change. However, should any incompatibilities exist, specific APIs that don't conform to the ABI definition of the new OS can be identified. Tools exist to support the static analysis of the application binary. Sun's `appcert` tool identifies differences so that they can be remediated prior to moving the application. This technology enables migration engineers to ascertain whether an application can be moved, without problem, to a newer version of the OS.

Moving applications from an older version of the Solaris OS to a newer version is referred to as *adoption* rather than migration. This distinction is made because the use of stable API standards, backward compatibility, and an ABI and tools enables you to verify and compare application interface usage, thereby guaranteeing that an application will run, without problem, in the new OS. In most cases, adoptions do not require recompilation, although using a later version of the compiler might provide performance benefits.

The E-Stack

Before we introduce migration and porting, consider the following figure, which we call the Enterprise stack or E-stack.

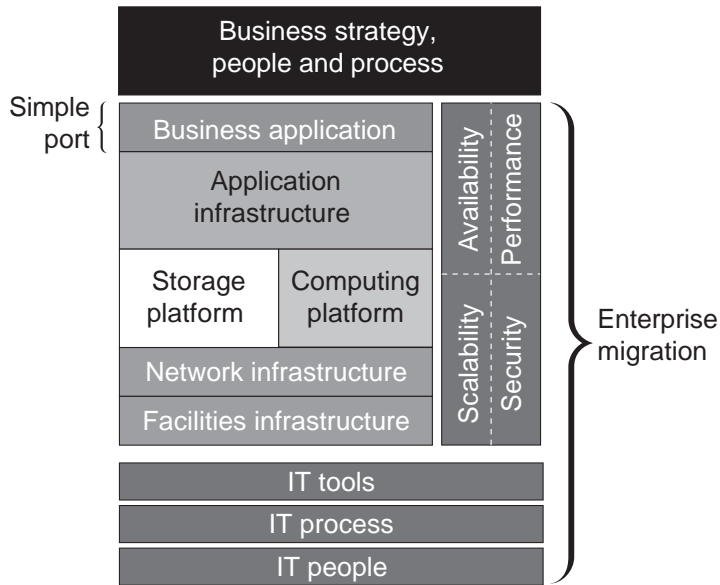


FIGURE 0-1 E-Stack Model for Enterprise Infrastructure

This construct is used as a model of the infrastructure of the enterprise. At the top of the stack, the business strategy, people, and process are defined. The high-level functions that occur here are usually controlled by an executive team. These functions provide logistical support to enable business functions, in addition to the unique value-added functions that differentiate one enterprise from another. The outputs of this level of the stack must be implemented by the lower levels. Typically, applications are used to implement and execute business processes.

The next section of the E-stack represents the execution architecture, which is composed of the following items:

- **Applications that support business functionality.** Changes in business processes most likely require changes in the application layer. Consequently, a rapidly evolving business must be able to implement application change in a timely manner. Business segments that are not subject to frequent change are less likely to require an agile development or runtime environment to effect application change.
- **Application infrastructure that supports the applications.** Modern implementations of this layer include web servers, application servers, middleware, and database technology that support a typical n-tiered architecture. For older, legacy applications, the application infrastructure layer is composed almost entirely of the OS. Applications in this layer are written to interact with the APIs provided by the OS vendor and software provided by independent software vendors (ISVs).

- **Compute and storage platforms.** These hardware components enable the application infrastructure. This layer of the stack is usually composed of a heterogeneous mix of hardware provided by a number of different vendors. As we will see, industry consolidation within the hardware manufacturing segment might require changes at this level of the stack. Unless the vendor has gone to great lengths to support backward compatibility, changes to the compute and storage platforms will require changes to the layers above them in the stack.
- **Network infrastructure.** In today's environment, the ability to communicate over a network is critical. This infrastructure can be based on any type of networking technology, from low-speed dial up to fiber-optic, high-speed backbones. Many legacy applications and their interfaces were designed and developed prior to the advent of networking technology. They depend on antiquated, proprietary interconnect technology, which might include a series of screens for data entry or possibly thick-client technology. These applications are not web enabled.
- **Facilities.** Frequently overlooked, this layer provides critical support to the stack elements above.

The upper portion of the E-stack that defines the execution architecture supports a number of systemic properties that are key to any enterprise. Availability, scalability, measurability, and security are all desirable, but can be implemented only to the extent that they are driven and supported by the execution architecture and its components.

The lower portion of the E-stack represents the management architecture. These tools, people, and processes implement the management infrastructure for the enterprise and combine to control, measure, and manage the execution architecture. Tools can be used to monitor capacity, utilization, and throughput, and to help ensure that service levels can be met. Processes are in place to support change, service, deployment, and maintenance. These tools and processes are selected, developed, and administered by the IT staff. As change is effected within the E-stack, IT staff must be made aware of all changes. Training must take place to ensure that people understand the management process, as well as the execution architecture.

As you can see, all the elements of the stack support each other. If the facilities do not provide adequate power or air conditioning, the results will manifest themselves in the compute and storage platforms. If the compute and storage platforms do not support the application infrastructure, the application will not be able to function correctly and service level agreements will not be met. This would mean that the business function and requirements mandated and defined from the top layer of the stack could not be implemented.

The following table outlines the relationship between consolidation, migration, and adoption.

TABLE 0-1 Consolidation, Migration, and Adoption

Term	Definition
Consolidation	The act of reducing the complexity in a data center
Migration	The act of moving technology from one platform or OS to another
Adoption	The act of moving from an earlier version of the Solaris OS to a later version

Porting

As illustrated in the E-stack shown in FIGURE 0-1 on page 4, the term *porting* applies to applications rather than infrastructures. In particular, it is usually used when talking about custom-written applications and refers to modifying or normalizing the code of an application so it can be recompiled and deployed on a new hardware platform that supports a different OS. Wherever possible, coding standards (ANSI, POSIX, and the like) should be adopted to minimize potential future changes that might have to be made.

Porting is inherently associated with modifying the code base of an application so that the functionality provided by the APIs of the existing OS and supporting software products is replicated in the new target environment. This is typically done by developing compatibility libraries that map the older APIs to the new environment. Vendors might provide these libraries to ease the burden of migrating applications to their environments, but in many cases, you will have to develop compatibility libraries yourself.

Porting the application requires minimal understanding of the logic or functionality of the application. It is a somewhat mechanical effort for making the application compatible with the new environment.

A porting strategy requires you to integrate the application with a new development environment, as well as a new operating system. While source code, scripts, and data are moved, compilers, source code repositories, and software tools are replaced by new versions that are compatible with the target platform.

When porting an application, you must also migrate any supporting third-party software. If the software is no longer available, you will have to find similar software and integrate it into the application. Should the amount of integration become excessive, the migration might begin to look less like a port and more like a

rearchitecture effort, as described later in the chapter. Ensure that you determine the availability of third-party software used by the application prior to choosing a migration strategy.

Migration

The term *enterprise migration* refers to migrating all layers of the E-stack, not only the application that supports business functionality. This is a very involved exercise that will have a greater impact on the entire IT staff than other strategies do. For example, migrations include the following changes:

- Management policies present in the old environment must be migrated to the new environment.
- Tools used to monitor and manage the execution environment must be replicated.
- Supporting software, in the form of third-party products provided by an ISV, or locally written scripts to manage applications and data, must be integrated into the new environment.
- People must be trained to administer the new hardware platform and operating environment.
- Adding a large symmetric multiprocessor compute platform might justify the use of a multithreaded architecture for the application.
- Implementing a Storage Area Network (SAN), rather than attached storage, might enable other applications to fully utilize storage resources that were previously unavailable.
- Adding networking capabilities might eliminate the need to transport and mount tapes.
- Web-enabling an application might reduce the need for proprietary terminal interfaces or thick-client technology.
- Changing hardware might require changes to the facilities to support more, or possibly less, power and cooling.
- Creating new tiers in the architecture might allow for the use of cheaper, more cost-effective hardware to support that portion of the application, which might, in turn, support greater availability and supportability.
- Using a modern programming language might enable the application to leverage more new third-party software, reducing the need for costly in-house development.

Evaluating the Environment

The migration solution you choose should be based on how an application fits into the overall IT environment. Consequently, you must evaluate the existing environment to determine both how the application meets business needs and its IT effectiveness.

Adequacy of meeting business needs relates to the application's ability to support the business functionality of the enterprise. Adequacy of meeting business needs can be defined as follows:

- The time required to introduce new features
- The ease of use of the application
- The ability to support the functional requirements of the enterprise
- The ability to support the future growth of the enterprise

Whereas the adequacy of meeting business needs relates to the application's ability to meet the current and projected functional needs of the enterprise, IT effectiveness measures the application's use of technology. IT effectiveness can be defined as:

- TCO
- Technological stability
- Functional separation
- Service level issues
- Implementation technologies

The comparison of an application's IT effectiveness with the adequacy of meeting business needs is represented in the following figure.

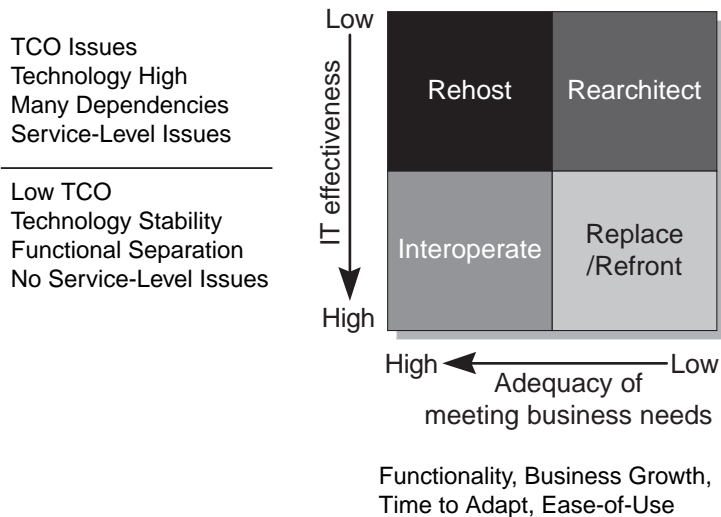


FIGURE 0-2 Effectiveness Versus Business Needs

The x axis evaluates how well the application currently fulfills its function in the business process. The y axis rates its IT effectiveness in terms of cost, technological stability, dependencies, and other factors. To determine a migration solution for each of the applications within the enterprise, we recommend that you begin by plotting the applications that have the greatest impact on the business process within this framework.

The evaluation of the application can be very formal (for example, a complete TCO study), or it can be ad hoc. Typically, the Chief Financial Officer of the enterprise will have to agree with the evaluation prior to agreeing to a budget to support the migration. By systematically evaluating all applications with the same criteria, a comprehensive migration strategy can be developed that will include a number of different solutions, as illustrated in the example shown in the following figure.

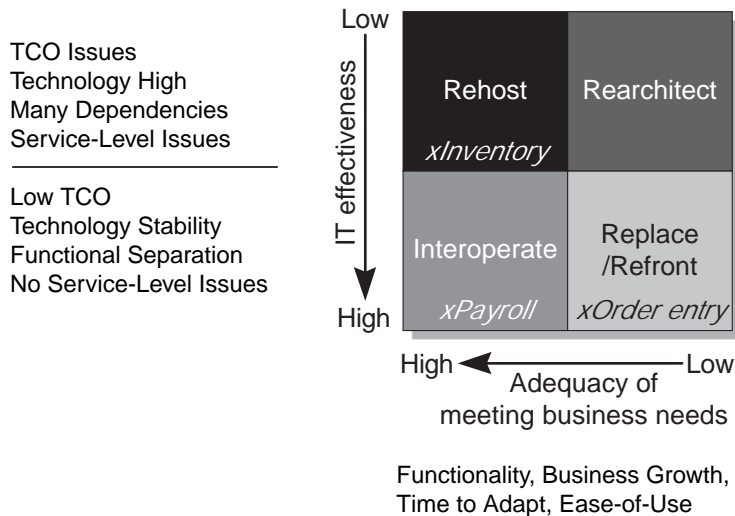


FIGURE 0-3 Example Evaluation of Applications

Applications with the following characteristics should be treated as follows:

- Applications that fall in the lower left quadrant are meeting business process needs and are highly effective with respect to IT; they probably should be maintained as they are.
- Applications in the lower right quadrant have high IT effectiveness, but are not meeting business process needs; they should be enhanced, not migrated.
- Applications that fall in the upper quadrants are the best candidates for migration. If they are meeting functional business process needs but are low in IT effectiveness (upper left quadrant), it is probably time to move them to a new environment. If they exhibit low IT effectiveness and lack functionality (upper right quadrant), it is probably time for the applications to be rearchitected.

FIGURE 0-3 illustrates how different applications (for example, payroll, inventory, and order entry applications) map to different migration solution spaces. For example:

- Order entry should be refronted because it relies on an antiquated user interface.
- The implementation and deployment of the payroll application is highly effective from an IT perspective and adequately meets business needs. It should remain where it is and be connected to new applications in the environment, as they are implemented, using connectors or adapters.
- The inventory application should be rehosted because the functionality it provides meets business needs, but there are issues with the hardware where it is implemented.

The terms “refront” and “rehost” are defined in detail in the following section.

Strategies

In the following sections, we examine the various strategies available for a migration effort, including the following.

- Refronting
- Replacement
- Rehosting—technology porting
- Rearchitecting—reverse engineering
- Interoperation
- Retirement

When planning a migration project, consider how your environment could benefit from the strategies described in this section.

Refronting

Many legacy applications have excellent functionality, but are not user friendly. Data entry for the application is accomplished by using a series of screens that frequently contain cryptic names for fields and unintuitive menus, which result from limited screen space. These interfaces were based on CRT technology that was available 20 to 30 years ago.

Rather than rewriting an entire application, it might be possible to change just the data entry portion of the application. *Refronting*, or adding a more aesthetic interface to an existing application without changing the functionality, is an option. Users will have access to the same data, but will be able to access it in a more efficient fashion without the use of expensive terminals, cabling, or peripheral interconnects.

Where desired and appropriate, a browser-based solution can be developed. In the case of mainframe replacement, 3270 data entry screens can be replicated over a network. Web-enabling an application can provide significant cost reduction. Different approaches for refronting include screen scraping, HTML generation, source code porting, and other techniques.

Modern graphical user interface (GUI) technology can also be integrated into a legacy application to support a clearer representation of the required input. Conversely, for reasons of efficiency, the data entry screens might be replicated in the new technology “as is” to eliminate the need to train the data entry staff.

Some of the user acceptance issues identified by this example might reveal themselves when a rehost strategy is adopted and a COTS product upgrade involves a change to the input form’s hosting technology (for example, when ASCII forms are replaced by a web browser).

The refront strategy requires an architectural model so that new components can invoke old, migrated components with minimum change to the migrated components. Such a migration project requires the application of architectural skills.

Replacement

The refronting strategy is really a variation of the much broader *replacement* strategy. Using the replace approach, the legacy application is decomposed into functional building blocks. Once it is broken down in this manner, portions of a generic and often complex, custom-written legacy application can be replaced with a COTS application. Of course, the package must be able to run on the target OS.

When evaluating replacement strategies, consider packages that offer better functionality and robustness than the existing, deployed application components. Make sure the vendor's solution is well tested and accepted in the marketplace, and verify that it is configurable, enhanceable, and well supported by the vendor. Product longevity and backward compatibility must also be taken into account.

One of the key drivers for the applicability of this strategy is the competitive dynamics of the software supply industry. Any custom (or bespoke) applications owned by users are always competing with the market, whether the competitive position is implicitly or explicitly evaluated. The marketplace is also driven by a sedimentation process. ISVs are seeking to maximize the value in business terms of their software products. Sedimentation refers to supporting functionality moving from the application implementation space to middleware (or utility software). From there, the functionality moves to the OS and often to hardware. For example, print spoolers and job schedulers are good examples of components that have been extracted from the application space and are now usually provided by utility software suppliers or by infrastructure vendors. Sometimes this moves even further, for example, in the case where web server load balancing functionality moved from the application layer to become an OS feature, and is now implemented in networking hardware.

The sedimentation process is an opportunity for migration planners because it enables their state-of-the-art business functionality to become available to the enterprise. This occurs because the ISV developers can outsource the functionality development and maintenance to alternative providers and can concentrate on transactional logic. By migrating some functionality through the use of the replace strategy, this trend can be copied, and the in-house code maintenance problem can be reduced, albeit through transfer to a third party. For example, cost can be reduced or developer wages can be more focused on benefits, but cost is not eliminated. The replace strategy enables in-house developers and maintainers of the utility code lines to be redeployed on more business-critical code lines and modules.

When considering replacement as a strategy, the option to replace the application's code with a new third-party software product might be attractive. If this approach is chosen, the migration project must do the following:

- Document the current business process and data model.
- Perform a gap analysis between the proposed application and the current state with respect to business process.
- Create a transformational data model, where appropriate.

These steps require traditional systems and business analysis skills.

A more powerful option might be to adopt a new application solution and change the organization's business logic when it no longer yields competitive advantage to match the package's optimum business process. This leaves migrators with the problem of identifying the still-used legacy data, and migrating to the new software solution. It also requires the development of a rollout plan that encompasses the enterprise's user community. Such a rollout is likely to be expensive, so the cost/benefit analysis of this approach needs to be solid and substantial. This involves data modeling skills and potentially programs for transforming the data into the target data model and populating the new database. The latter approach allows the enterprise to transform applications built to deliver functional competitive advantage to software built to allow the user organization to compete through superior cost advantage. This mandates that the replacement product be competitively inexpensive to deploy and run.

Replacement can be a quick, low-risk solution, although the replacement of complete applications will have large implications in terms of business acceptance and rollout. However, replacing a homegrown ERP solution with a COTS package can also take upwards of two years. Effort is required to ensure that business processes and logic now conform to the capabilities of the COTS component, rather than the other way around. For some applications, the cost of acquiring custom logic for these software packages can be equivalent to maintaining and modifying a custom code base, depending on the function provided by that package, which is why it might be more appropriate to adopt the COTS vendors' assumed business process. Not all business processes, and hence not all applications, are designed to enable functional competitive advantage. For instance, a customer relationship management (CRM) package deployed to replace a specific business function will require more maintenance than configuring a replacement print spooler. If the proposed source modules for the migration are only a subset of the target package's functionality, it might make more sense to identify additional business processes to encapsulate within the CRM solution. For example, you might replace more code and increase the potential benefits case. This shows the trade-offs available when planning replacement strategy-based migrations.

There are three clear options within the replacement strategy.

- Use a COTS package to replace or retire the source modules.
- Use a COTS/utility package to replace sedimented functionality.

- Use operating system functionality to replace sedimented application functionality.

The last option in the preceding list uses functionality that has been integrated into the existing OS. Examples of this include complicated memory management schemes that have been implemented due to older memory limitations, coarse-grained parallelism that is used instead of threading models, or shared memory that is used as an improved IPC shared memory.

The advantages of moving from a homegrown solution to a COTS-based solution include:

- Integration with other internal and third-party external applications
- The release of the budget associated with inflexible development resources
- An improved opportunity to tap skilled resources from established labor markets supporting both the business and IT communities

Replacement can act as a strategy on its own, and it can also be applied to components within an alternative strategy. Interestingly, as a strategy, it potentially yields the highest benefits and involves the highest degree of cost, yet when applied to components within an alternative strategy, it can be a quick and low-risk strategy.

Rehosting—Technology Porting

Rehosting involves moving complete applications from a legacy environment with no change in functionality. There are several ways this can be accomplished for custom-written applications.

- **Recompilation.** As previously mentioned, an application can be ported to the new environment. There are two approaches for doing this. The first approach is primarily associated with developing or acquiring a compatibility library that provides identical functionality to that of the APIs found on the original OS and supports third-party products. For example, Sun provides compatibility libraries for some of the major competing operating systems such as HP/UX. An alternative approach is to use intelligent code transformation tools to alter the original source code to correctly call the new operating system's APIs. Both of these approaches have the benefit of capturing the changes required during the migration, although the second might limit backward compatibility.
- **Emulation.** This approach introduces an additional software layer to emulate the instruction set used in the source binaries. While introducing another software layer between the application and the hardware can affect performance, it eliminates the need for recompilation. When adopting this strategy, it is important to understand that the old environment has not really been left behind. The application will be developed and compiled using the old environment and will

only execute in the new environment. By their nature, emulation solutions incur additional cost above that of the target platform environment. This results from the need to supplement the OS with the emulator, which is rarely free.

Although emulation is a useful approach, if source code is available, it is more common for an application to be recompiled to the native instruction set because native code runs faster. Emulation is most useful when migrating applications that are written in interpreted languages or when the original execution environment was tightly coupled within the OS. BASIC, PICK, or MUMPS are examples of development or runtime environments that are suitable for emulation solutions.

Most emulators are for interpreted languages; therefore, the source code is available to the organization. However, source code engineering and reverse engineering rights might not have been granted in the right-to-use license. If you intend to use an emulation solution or reverse engineer a solution, ensure that you are licensed to do so in the environment where it will be used.

- **Technology porting.** Emulation is a technique that supplements the target environment with the capability to execute code (usually interpreted) that runs natively on the original system. Many applications are developed and written in a superstructure software environment that is installed as a layered product on the source system. The most common of types of these applications are created by relational database management system (RDBMS) vendors, many of whom support an array of hardware platforms and guarantee a common API across those platforms. The advantage of this approach is that one common API owner, the software ISV, owns the API on both the source and target systems. While the discovery stages of a migration project are still required, the APIs on the source and target systems remain the same.

The leveraging of the ISV solution is often an opportunity to upgrade the ISV product version to obtain new functionality or to obtain superior support from the ISV. For instance, the transaction processing system known as CICS relies on a well-understood series of APIs. These APIs and their functionality have been ported or reimplemented on the new target Solaris OS. Applications using these APIs are compiled to run native instructions on the new system.

Rehosting offers the advantage of low development risk and enables familiar legacy applications to be quickly transferred to a more cost-effective platform that exhibits lower TCO and a faster return on investment (ROI). Extensive retraining of users is not needed because the architecture, interface, and functionality do not change. Rehosting is an excellent approach for companies desiring to decrease their maintenance and support costs.

Rehosting is, by definition, a quick fix. Rehosting does not change the application or the architecture. This means that new technology that is available in the target environment might not be properly utilized without some modification of the application. Rehosting is a preferred solution when the current business logic and business process remain competitive in the enterprise's markets and are worth

preserving. Rehosting offers the possibility of using cost savings accrued through switching development and runtime environments to fund full rearchitecture projects, where warranted.

Rearchitecting—Reverse Engineering

Rearchitecting is a tailored approach that enables the entire application architecture to migrate to the new OS, possibly using new programming paradigms and languages. Using this approach, applications are developed from scratch on a new platform, enabling organizations to significantly improve functionality, and thereby take full advantage of the full potential of a target system.

Applications poor in IT effectiveness and functionality are the best candidates for rearchitecting. This approach is best utilized when time is not a major factor in the decision. Most rearchitecture projects require a skilled development staff that is well versed in the new technology to be implemented.

The downside to this approach is that it requires new or additional training for users, developers, and technical staff. In addition, rearchitecting requires the most time and is the most error prone of all of the possible solutions. Sometimes, business rules can be well hidden in user interface or database management systems. For example, this was the case with DECforms, DEC FMS, and any RDBMS triggers. The ability to extract all the necessary business logic from the application's source can be severely inhibited by poor coding methodology and practice. An example is the hard coding of business parameters.

Despite these problems, it remains that re-architecture and reverse engineering are perceived to be the correct strategies, and these problems become project risks. These risks can be mitigated by the application of appropriate business acceptance testing with internal and external users.

Rearchitecting does, however, open the opportunity to improve the business logic and processes and to change the developer productivity model.

A technique particularly appropriate to rearchitecture is reverse engineering. It is an axiom that the business logic encapsulated in the source code is the business logic implemented, and thus the source code is the most accurate place to discover the business logic. One of the key problems of software development is that most usability errors in software are introduced by poor business process documentation and even poorer translation into software idioms. Some software environments have embedded dictionary or repository functionality. Where these exist, they may be supplemented with original author or third-party tools to enable the extraction of business logic and the recreation of that business logic in new environments. With these tools, the process can be reversed, the dictionary can be parsed, the user world view can be generated, and the implementation source code can be generated. A classic example is the RDBMS world where data definition language scripts for a

database implementation can be generated from the database implementation itself. This is tool based, and tools might be proprietary to a single RDBMS or environment, or they might be open, running across multiple environments. Database schema generators are particularly useful for migrating from one RDBMS to another, such as from Microsoft's SQL Server to Oracle or Sybase.

Interoperation

In certain cases, it might be advantageous to leave an application where it is and surround it with new technology when it is required by an enterprise. *Interoperability* is a strategy that should be considered in the following cases:

- If business requirements are being met and IT effectiveness is high, it might be desirable to leave the application in its current environment, provided that environment is capable of interacting with current technology.
- Unfortunately, business drivers—for example, the existence of a leasing or outsourcing contract—might dictate that an application should stay where it is for some period of time. This is one of the risks of abandoning your IT environment to a third party. Over time, outsourced applications become orphans within the IT infrastructure. They are not fully integrated into the IT environment, and they most likely do not have a development staff and run on outdated hardware that is no longer cost effective.

Many ISVs currently provide technology that enables legacy applications and storage technology to interoperate with newer technology. Intelligent adapters exist that support interactions between the mainframe and more modern computing alternatives. It is also possible to compile an older language such as COBOL or PL/1 into Java™ byte-code, enabling it to seamlessly interact with a modern application server and other components of a Java™ 2 Enterprise Edition (J2EE) environment.

When choosing this strategy, it is important to understand the vendor's commitment to the existing product line, as well as any future maintenance and product licensing costs. In addition, consider the availability of third-party software and current technological trends. Where possible, open standards should be favored to allow a wide choice of competitive options.

Retirement

Changes in technology can obviate the need for specific functionality in an application or an overall solution. As middleware or third-party products mature, they might render the functionality implemented in the application obsolete. In this case, legacy utilities or legacy application functionalities can be retired because they are no longer required or are implemented elsewhere in the solution.

Choosing a Strategy and Developing Tactics

The ideal migration solution incorporates a number of the strategies listed in this chapter, where appropriate. Each of the strategies identified in the preceding section has a number of closely aligned supporting techniques. The selection of a strategy will define the obvious and most effective technique, but it might need to be supplemented with techniques more appropriate to other strategies.

The following table summarizes the strong alignments of particular migration techniques with the migration strategies.

TABLE 0-2 Strategy and Technique Alignment

Strategy	Complementary Technique
Refroniting	Redeveloping, reverse engineering, source code porting
Replacement	Reverse engineering
Rehosting	Source code porting, technology porting
Rearchitecting	Reverse engineering, redeveloping
Interoperation	Technology porting, emulation
Retirement	Reverse engineering

As with the migration decision itself, the tactical approach used to solve a problem or to provide functionality must be evaluated in terms of its effectiveness, its impact on the ability to meet business needs, and its cost. Tactical decisions made to resolve technical issues might impact the overall project, either beneficially or adversely, in terms of systemic qualities, manageability, training, and cost.

As described in the preceding sections, there are a number of migration solutions for you to choose from. Each has its own benefits, as well as its own drawbacks. Selecting the correct solution should realize the associated value proposition.

One factor in your decision should be based on your application architecture. Most modern application architectures are now based on n-tier models. This decomposition allows for different strategies to be applied to different tiers, where appropriate. This might mean that more than one strategy will drive your migration.

Conversely, older legacy applications might be monolithic or client-server in design and implementation, which offers the opportunity to rearchitect to an n-tier model. The following table outlines some common (not exclusive) approaches for each tier.

TABLE 0-3 n-Tier Migration Strategies

Tier	Purpose	Common Approaches
Presentation	Hosts the processing that adapts the display and interaction as appropriate for the accessing client device, be it a desk desktop computer, a cell phone, a PDA, or any other device.	Refronting, rehosting, interoperating, and replacing
Application or Business Logic	Hosts the logic that embodies the rules of the enterprise, irrespective of access device or resource implementation.	Rehosting, interoperating, and replacing
Integration	Allows for the connection of disparate applications and data sources.	Rehosting, interoperating, and replacing
Resource or Database	Consists of legacy systems, relational databases, data warehouses, or any other back-end or external processing system which accesses and organizes data.	Rehosting and replacing
Persistence	Holds the permanent data for the enterprise. In the past, this was considered part of the Resources tier, but with the growth of intelligent storage (SANs, NAS and intelligent arrays) it has become a tier in itself.	Rehosting and replacing

Another factor to consider is the relationship between value and effort, as shown in the following figure. Typically, value is proportional to the amount of effort that is expended on a project. In the following paragraphs, we examine each of the proposed migration solutions as they relate to value versus effort.

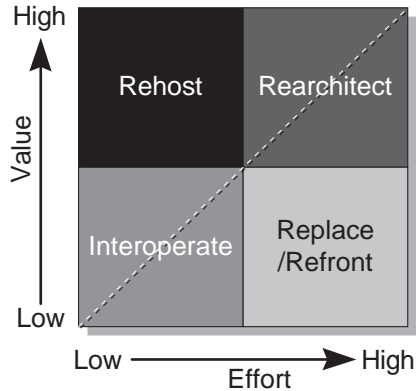


FIGURE 0-4 Relationship of Value to Effort

- **Interoperation.** This solution requires the least amount of effort, but also provides the least amount of benefit. The existing architecture and infrastructure stay in place and simple connector technology is deployed to support the interaction with new applications or hardware that might be deployed. Because no new functionality is introduced, this effort requires minimal time and expense.
- **Rearchitecting.** This solution occupies the other end of the scale. Rearchitecting the application has great benefits, because it supports tailored functionality; modular, tiered design; and a modern implementation language. However, the amount of effort (and associated cost) can be significant. As well as incurring significant expense, time, and effort, this solution can also introduce errors, so it requires a rigorous validation and verification effort.
- **Refronting or replacement.** This solution is targeted towards enhancing an application that has already been deemed to be somewhat unacceptable in meeting business needs. This solution is targeted to applications that are not meeting business needs, but that do not have any IT-related problems relating to Service Level Agreements (SLAs), QoS, TCO, and the like. Enhancing the application by adding a presentation layer will add new functionality, but given that the application has already been found to be somewhat unacceptable, this enhancement adds minimal overall benefit compared to the amount of effort that it requires.
- **Rehosting.** As illustrated in FIGURE 0-4, rehosting is the solution that provides the most value for the least effort. Rehosting typically involves modifying the source code and build environment for an application so that it compiles and runs on the new target system. During this process, new features and functionality are not added. Many companies often want to add new features or functionality when they migrate, but these steps should take place after the application has been migrated.

Business logic remains the same when an application is rehosted. The only application changes usually relate to the APIs of the target OS. Over time, standards (SVR4, POSIX, and the like) have converged so that differences between versions of UNIX are minimal; therefore, migrations of applications between different versions of UNIX require minimal effort.

Rehosting applications from proprietary, non-UNIX environments that do not adhere to open standards can prove to be more challenging.

In certain cases, an application must be changed to not only adhere to the APIs of the OS, but to interact with third-party product code as well. For example, consider the rehosting of a CICS application from MVS to MTP/MBM application running on the Solaris OS. The application's interaction with the MVS environment must be recoded in such a way that similar calls and functionality are used in the Solaris OS, but the CICS interaction requires minimal conversion because the CICS functionality and APIs have been redeveloped under the Solaris OS by the MTP and MBM product set.

Rehosting has the following characteristics:

- **Least expensive and requires the least effort.** Rehosting requires minimal changes to be made to applications for them to run under the new environment. Therefore, the cost and effort involved in this strategy are minimal.
- **Quickest implementation.** Because little or no code is written and no new functionality is added, this solution can be completed in minimal time, compared to the other migration solutions.
- **Business logic remains the same.** Rehosting doesn't typically include the addition of new features or functionality. Consequently, the business logic remains the same, meaning that minimal or no staff training is required and few, if any, changes have to be made to the organizational structure.

Case Studies

In the last three chapters of this book, we examine three case studies, each of which uses a different migration strategy.

- **Case 1: Small business, Linux.** This example is based on a small software and services development company that is looking to move from Linux to the Solaris OS. There are approximately 20 servers being used: 10 for production, 5 for development and testing, and 5 for office support tasks. Their application is mostly Java-based, but they use MySQL for the database. Significant shell scripting has also been used for utilities in the product.

- **Case 2: Custom application, Tru64.** In this example, we examine the migration of a mythical inventory application implemented in the C programming language. The application is integrated with a Sybase database running under the Tru64 environment. The exercise involves porting the application so that it runs under the Solaris environment and replacing the Sybase database with an Oracle relational database.
- **Case 3: General ledger, HP-UX.** In this example, an insurance company planned to move its accounting, risks, and claims software from HP/UX to the Solaris OS to achieve superior scalability against planned business and system growth. This exercise utilizes the rehost strategy, in which the technology porting approach is used to minimize risk and cost.

Related Resources

This article is an excerpt from the Sun BluePrints book *Migrating to the Solaris Operating System*. Refer to the book for more information about the topics presented in this article.

About the Authors

Ken Pepple is an IT Architect in the Sun Professional Services (SunPS) Asia Pacific practice. In this role, he assists clients with enterprise computing architectures, concentrating on advanced data center projects. In addition to these activities, Ken recently co-authored, with David Hornby, the Sun BluePrints[tm] book, "Consolidation in the Data Center: Simplifying IT Environments to Reduce Total Cost of Ownership".

Before moving to his current position, Ken managed the SunPS high-end platform services program, and focused on complex performance issues for the IT Consulting and Operations practice. While there, he co-authored and taught the Sun Education seminar "Solaris Performance and Tuning Secrets".

Brian Down has been at Sun for five years in Sun's Global Sales Organization as a Senior Staff Engineer and most recently in SunPS where he holds the position of Chief Architect for Enterprise Migration for the Americas. For several years, Brian has focused on developing Sun's migration methodology and solution strategy, helping to develop and identify the value proposition associated with such implementations. Prior to joining SunPS, Brian focused on performance and custom engineering initiatives related to strategic server installations for the GSO.

With over 25 years of industry experience, Brian has held various engineering positions ranging from senior engineer with a computer security company to that of Research Associate for the Department of Computer Science at the University of Toronto, where he worked for over 10 years.

David Levy is a Senior Consultant in the SunPS United Kingdom (UK) organization. He is currently leading the program for the UK's Consolidation and Migration team, concentrating on data center architectures and economics. Prior to this role, Dave led the UK's financial services consulting team based in London. Dave has successfully completed numerous consolidation and migration projects for banking and media customers.

Before working for Sun, Dave worked for a number of financial services, IT manufacturing, and government organizations, primarily as a database architect and engineer. Dave has authored presentations for the Oracle and Sybase user groups and is a member of the British Computer Society and Chartered Institute of Management. Dave is an Honors graduate of the University of Exeter, where he majored in Economics.

Ordering Sun Documents

The SunDocsSM program provides more than 250 manuals from Sun Microsystems, Inc. If you live in the United States, Canada, Europe, or Japan, you can purchase documentation sets or individual manuals through this program.

Accessing Sun Documentation Online

The `docs.sun.com` web site enables you to access Sun technical documentation online. You can browse the `docs.sun.com` archive or search for a specific book title or subject. The URL is `http://docs.sun.com/`

To reference Sun BluePrints OnLine articles, visit the Sun BluePrints OnLine Web site at: `http://www.sun.com/blueprints/online.html`