

OPEN PLATFORM MIGRATION GUIDE 2018

TABLE OF CONTENTS

DEVELOPING A MIGRATION PATH TO OPTIMIZE I.T.	2
BENEFITS OF MIGRATING.....	2
MIGRATION PATH.....	2
PAVING THE WAY TO AGILE I.T.....	3
OVERVIEW: MIGRATING TO X86 ARCHITECTURE	4
ASSESSING THE APPLICATIONS TO BE MOVED TO THE NEW INFRASTRUCTURE.....	4
ASSESSING THE LOAD ON CURRENT SERVERS	5
IDENTIFYING COMPELLING PILOT OPPORTUNITIES.....	5
DEVELOPING A PROOF-OF-CONCEPT APPROACH	5
RE-EVALUATING THE SOLUTION ARCHITECTURE	6
PERFORMING A REHEARSAL MIGRATION.....	7
LAUNCHING THE PRODUCTION DEPLOYMENT	7
NEXT STEPS.....	7
MIGRATING APPLICATIONS FROM UNIX TO LINUX	7
THE MIGRATION SPECTRUM.....	7
MIGRATION CONSIDERATIONS	8
REASONS FOR MIGRATING	9
COMMON MIGRATION SCENARIOS	9
MIGRATION DEPLOYMENT SCENARIOS	11
CONSOLIDATION	12
DISPERSION.....	13
BEST PRACTICES FOR A SUCCESSFUL MIGRATION	14
GENERAL GUIDELINES FOR MIGRATING ENVIRONMENTS.....	14
BUILDING A STANDARD OPERATING ENVIRONMENT	15
PROVISIONING COMPONENTS.....	16
OVERVIEW OF RED HAT SATELLITE	17
CREATING A STRATEGIC PLANNING ROADMAP	18
CONSOLIDATED ANALYSIS OF EXISTING HARDWARE.....	18
CONSOLIDATED DEPLOYMENT SCENARIO AND VIRTUALIZATION ANALYSIS.....	18
HIGH-LEVEL HARDWARE REDEPLOYMENT ANALYSIS.....	19
CONSOLIDATED RISK ANALYSIS AND RISK MITIGATION PLAN UPDATE.....	20
TRAINING PLAN CREATION.....	20
DEEP-DIVE ANALYSIS OF LARGE, HIGH-COMPLEXITY APPLICATIONS (OPTIONAL)	20
COST ESTIMATION.....	20
CREATING A MASTER MIGRATION ROADMAP.....	21
SELECTING THE OPTIMAL INTEL® XEON® PROCESSOR-BASED SERVERS.....	21
CONTINUING THE RED HAT I.T. MODERNIZATION JOURNEY	22
RED HAT MIGRATION SERVICES	23
INTEL RESOURCES AND SERVICES	23

IT optimization represents a first step in a broader move toward hyperconvergence that integrates compute, storage, networking, and virtualization resources within a COTS framework and software-centric architecture.

Over many years, Intel and Red Hat have worked collaboratively, optimizing Intel architecture platforms running Red Hat Enterprise Linux to perform reliably, securely, and effectively.



facebook.com/redhatinc
[@redhatnews](https://twitter.com/redhatnews)
linkedin.com/company/red-hat

DEVELOPING A MIGRATION PATH TO OPTIMIZE I.T.

Businesses thrive on technology innovation, which serves as a virtual engine to generate value and spur growth. Costly, inefficient traditional IT infrastructures present barriers to innovation, create ongoing maintenance difficulties, and introduce serious end-of-life concerns as both hardware and software components become obsolete.

BENEFITS OF MIGRATING

While traditional IT infrastructures may be necessary in the near-term as enterprises work out strategies for long-term migration and modernization, businesses can benefit from taking these migration steps:

- Moving from a proprietary operating system (OS) to Red Hat® Enterprise Linux®.
- Adopting a foundation based on Intel® architecture-based hardware.

Business values strengthened by following this basic migration path are quickly evident, both in terms of capital expenditure (CAPEX) and operating expense (OPEX) savings. By dropping expensive proprietary licensing agreements in favor of a predictable, cost-effective subscription model offered by Red Hat, OPEX is reduced substantially. CAPEX is lowered as well by replacing pricey proprietary systems with commercial off-the-shelf (COTS) servers powered by Intel, delivering scalable, high-performance compute power with the capacity for greater, virtual machine density, making maximum use of IT resources. Combined with Red Hat® Virtualization, Intel® processor-based servers represent an innovative path to more agile and elastic IT resource availability, at a much lower cost than proprietary alternatives.

The standards-based hardware and software framework coengineered by Red Hat and Intel also provides access to a vast ecosystem of community-powered solutions that have been tested and verified on reference architectures. Enterprises planning to advance modernization efforts further can take advantage of Red Hat's scalable management platforms and interoperable solutions. Staff efficiency can be increased significantly with better server-to-admin ratios and access to a centralized management console across physical, virtual, and cloud workloads. IT optimization represents a first step in a broader move toward hyperconvergence that integrates compute, storage, networking, and virtualization resources within a COTS framework and software-centric architecture.

MIGRATION PATH

The general migration path discussed in this guide consists of four basic steps:

1. **Planning.** Determine the workloads that you want to move from the proprietary system to the Intel and Red Hat open platform.
2. **Converting.** Perform the necessary code porting and data conversion in preparation for the migration, and determine if additional third-party libraries or middleware need to be available in the new environment.

3. **Configuring hardware:** Determine the Intel architecture-based hardware configuration that is best suited for supporting the anticipated workloads.

4. **Deploying:** Complete the deployment to the new hardware in a staged series of steps. Most deployments require these migration stages:

- Proof of concept (PoC)
- Pilot migration
- Rehearsal
- Production launch

The actual steps required, however, depend on the environment being rehosted, which vary according to individual circumstances. Figure 1 shows the fundamental steps as they apply to typical applications involved in a migration initiative.

Environment	1. Project planning	2. Code conversion	3. Proof of concept	4. Solution architecture	5. Pilot migration	6. Rehearsal migration	7. Production migration
Infrastructure	★				★		★
Remote office/retail computing	★	★		★	★	★	★
Business-critical COTS applications	★	★	★	★		★	★
Business-critical custom applications	★	★	★	★		★	★

Figure 1. Migration methodology and steps

PAVING THE WAY TO AGILE I.T.

In the migration from traditional, proprietary IT environments to modern, agile platforms, the choices at each stage of the process can be overwhelming. Many IT directors, CTOs, and IT staff members recognize the prospective business benefits that can be achieved by moving from UNIX* to Linux (U2L) and from a reduced instruction set computer (RISC) architecture to standard high-volume servers (SHVS) based on Intel architecture hosting an open source software stack. But the individual steps involved in moving from a proprietary platform to an open platform can be complex and daunting. Careful planning and expert guidance can help avoid potential pitfalls and better align with proven best practices for this type of migration.

“A five-year ROI analysis completed by IDC in August 2016 determined that, for the Red Hat Enterprise Linux and Red Hat Satellite customers surveyed, the ROI would reach 373 percent and customers would break even on their investments, on average, in five months.”

AL GILLEN, MATTHEW MARDEN,
SEPTEMBER 2016³

Analyst firms have approached the process of strategically upgrading IT infrastructures using different terminology but with essentially the same concepts. IDC Gartner uses the term bimodal¹ to describe two different modes of IT delivery: one engineered for stability and the other for agility. IDC looks at the progressive upgrading in terms of a 1st, 2nd, and 3rd platform,² where the 3rd platform is based on full digital transformation and cloud-native applications.

The important point is that enterprises do not need to launch a digital transformation effort that immediately moves the entire business into the cloud. The first measured step for many businesses is an on-premise upgrade of the servers running business operations to SHVS and the OS environment to Red Hat Enterprise Linux. Business benefits accrue from the moment this optimized infrastructure is put into operation.

The analyses of total cost of ownership (TCO) and return on investment (ROI) for moving from both RISC to Intel architecture and UNIX to Red Hat Enterprise Linux have been developed by Alinean, supporting the cost benefits associated with migration. These findings are presented later in this document.

For additional insights into the joint Intel and Red Hat value in optimizing and modernizing an organization's datacenter, download the whitepaper [“Alleviate technical debt.”](#) More about the ways in which Intel and Red Hat work together on innovative IT solutions can be found on the [strategic alliance web page](#).

OVERVIEW: MIGRATING TO X86 ARCHITECTURE

Accomplishing a successful migration to an x86-based infrastructure requires a careful, methodical approach based on best practices. Every situation is unique, but the following steps provide an overview of the most important considerations and recommended approach for IT managers planning the deployment. This planning process is discussed in more detail later in the guide.

ASSESSING THE APPLICATIONS TO BE MOVED TO THE NEW INFRASTRUCTURE

Take time to assess the current UNIX environment and identify equivalent capabilities provided by the Red Hat Enterprise Linux ecosystem. This assessment should include third-party business applications and determine the equivalent capabilities in the Red Hat Enterprise Linux ecosystem, as well as any risk factor associated with the migration. At this stage, develop a strategic migration plan to Red Hat Enterprise Linux that features a comprehensive roadmap and cost assessment.

Red Hat recommends following this well-tested process that identifies migration opportunities, examines the risks associated with various migration scenarios, creates a standard enterprise build, and develops a comprehensive strategic migration plan for the enterprise:

- Examine the existing UNIX architecture and determine the equivalent capabilities in the Red Hat Enterprise Linux ecosystem.
- Examine third-party functional and business applications and determine the equivalent capabilities in the Red Hat Enterprise Linux ecosystem.
- Measure organizational readiness and overall migration risk.

¹ “Gartner IT Glossary.” Gartner. 2017. <https://www.gartner.com/it-glossary/bimodal/>

² “The Shift to the 3rd Platform.” IDC. 2017. <https://www.redhat.com/cms/managed-files/cl-shift-to-third-platform-idc-analyst-infographic-f9614-201710-en.pdf>

³ “The Value of Standardizing on Red Hat Infrastructure Solutions.” IDC. 2016. <https://www.redhat.com/en/resources/value-of-standardizing-red-hat-infrastructure-solutions>

- Develop a strategic UNIX to Red Hat Enterprise Linux migration plan, including a detailed roadmap and cost estimate.
- Implement the strategic migration plan and employ implementation support strategies.

ASSESSING THE LOAD ON CURRENT SERVERS

Once the selected application migration has been mapped out, perform a thorough load assessment on the servers so you can determine an appropriate configuration for PoC testing. Preliminary server sizing helps determine an appropriate configuration for the workloads that will be handled. Application performance can differ substantially when moving from one architecture to another. Performance testing on similar platforms ensures an accurate performance estimate. At this stage, you only need a reasonable preliminary estimate to support your test requirements—the actual PoC provides you with the necessary details to size the production server.

Determining memory and disk space for the test server can be accomplished in a fairly straightforward manner, but the implementation is likely to be quite different than the original system because of the availability of greater compute resources and expanded functional capabilities. Selecting the processor best suited for the test server is another key consideration. Determine the appropriate numbers of cores, threads, cache, instructions per clock cycle, and other processor specifications, keeping in mind the anticipated workload once the platform migration is completed. Performance-monitoring tools available in most UNIX operating systems can help with this task.

IDENTIFYING COMPELLING PILOT OPPORTUNITIES

The pilot establishes the essential parameters for building an operating environment that effectively supports the migration effort. Identify the best initial application to migrate based on current business requirements, ideally positioning the enterprise for an early win. Even though it may be tempting to select a simple and obvious application for the initial migration, this approach may not provide the strategic traction needed to convince business units of the importance of continued migration efforts. A business unit that has a strong history of adopting modern technology and embracing change—as well as a willingness to share its IT successes across the organization—represents a good target for a migration.

DEVELOPING A PROOF-OF-CONCEPT APPROACH

All complex or business-critical migrations should include a PoC to verify that the application can run in the new environment and determine how it performs under anticipated workloads. The PoC can also help determine the best ways to optimize configurations for maximum performance and establish the preliminary process for production migration.

To conduct the PoC:

- Ensure that you have made the necessary provisions for migration testing.
- Verify that staff members from the appropriate business unit are available to oversee the testing.
- Confirm the physical location of the server(s) and the database that will be used for test migrations.
- Establish provisions for testing the application's performance on the target hardware.
- Ensure that you lock down any changes to the application during migration.

At this stage, completing a more thorough evaluation for the server sizing and application scaling can help improve the success of the PoC. Factors that can impact server platform selection and sizing include:

- Initial measured average utilization.
- Maximum utilization target.
- Peaks in demand.
- Workload growth projections.
- Relative capacities of two-socket and four-socket servers.
- Consideration of necessary application scaling to conform with service-level agreement (SLA) requirements.
- Advanced considerations, including clustering, disaster recovery, and failover.

More information about conducting a PoC can be found in the paper [“Migration from UNIX/RISC and Mainframe to Intel-based Solutions.”](#)

RE-EVALUATING THE SOLUTION ARCHITECTURE

After completing the PoC, you will have gained assurance that the selected applications to be migrated to the new environment can run efficiently on the hardware and that SLAs can be met. This is a good point at which to re-evaluate broader issues, including the scalability and availability of the solution. Key issues to be evaluated include:

Robust hardware. With most migrations, you automatically gain performance and scalability, as well as reliability, availability, and serviceability (RAS) when you replace older equipment with newer, more robust hardware. You can add to these advantages by configuring new systems with redundant components, such as power supplies, disk drives, and fans. For a high level of reliability and top performance for business-critical applications, the Intel® Xeon® processor Scalable family features advanced RAS capabilities and exceptionally high availability. Processors in this family also deliver high levels of compute density and can scale to accommodate massive workloads. Built-in processor architecture enhancements provide exceptional handling of high-performance computing (HPC) workloads.

Clustering. To reduce the effects of failure within the migrated environment, use high-availability clustering software. This software is available through Red Hat Enterprise Linux 7 High Availability Add-On. This add-on features an integrated set of components that support different configurations for determining performance, high-availability, load balancing, scalability, file sharing, and cost effectiveness. Clustering designates a redundant server to take over the duties of the main server in the event of failure. Failover is not instantaneous, but it is automatic. Clustering software limits the effects of failure on overall system availability.

Horizontal scaling. Horizontal scaling handles increased workloads by adding servers rather than increasing the size of individual servers. This approach works well for presentation layer services, which tend to scale well across multiple servers through workload balancing. If a server fails, its workload automatically fails over to the remaining servers, providing high availability as well as flexible scalability. Adding automated provisioning and maintenance tools keeps overhead management low as the solution expands.

Vertical scaling. Applications within a vertically scaled architecture reside on a single server. As workloads grow, performance is scaled by adding resources to the server or upgrading to a larger system. Consider mature or single-purpose applications that require a high degree of scalability for use with vertical scaling.

PERFORMING A REHEARSAL MIGRATION

Ensuring a successful switchover to the new solution requires careful preparation. A rehearsal migration makes it possible to improve the process used in the pilot migration and streamline it for maximum efficiency. By evaluating and reordering your process, developing scripts to automate necessary processes, and carefully examining data migration steps, you can gain knowledge and experience to help ensure a smooth production migration. Document the steps involved, and then test the solution for quality assurance (QA) and acceptance. More importantly, you should be prepared to repeat the rehearsal migration as necessary until you are confident about proceeding.

LAUNCHING THE PRODUCTION DEPLOYMENT

At this point, you should have full confidence in the new solution and the migration process.

- Use your rehearsal documentation to develop a tight project plan.
- Schedule an optimal maintenance window and notify all affected business units.
- Rebuild the target server and finalize all premigration steps.
- Execute the production migration as rehearsed.
- Conduct QA and acceptance.
- Cut over or run in parallel until you receive sign-off.
- Document each step that was completed.

NEXT STEPS

The following sections provide a more detailed discussion of the steps involved in moving core applications from a UNIX environment to one based on Red Hat Enterprise Linux. The most common migration scenarios in an IT optimization are described, as well as the level of difficulty associated with specific types of application migration.

MIGRATING APPLICATIONS FROM UNIX TO LINUX

The scenarios described in this and the following section address migration from a generalized UNIX architecture—that might consist of AIX*, Solaris*, and HP-UX* environments—to a Red Enterprise Linux ecosystem. Additional documentation and consultation is available from Red Hat to support a full range of migration scenarios.

THE MIGRATION SPECTRUM

A successful migration effort must take into account the types of applications to be moved to the new platform, the relative difficulty in converting code and data, and the hardware platform best suited to host the planned Linux environment. The relative ease of different application migrations is shown in Figure 2.



	Infrastructure applications	Remote office/retail computing applications	Business-critical COTS applications	Business-critical custom applications
Code/data conversion	Little or no code conversion required	Complexity varies depending on the applications	Low complexity	Moderate complexity
Downtime/outage	Not significant	Can be controlled or scheduled	Downtime sensitive	Downtime sensitive
Replication across environment	Easy to replicate	Pilot then replicate	Varies from easy to moderate, depending on the apps involved	Varies from moderate to complex, depending on the apps involved
Cost and risk	High value with low cost and risk	Cost and risk are defrayed	Moderate risk and cost, depending on recoding complexity or replacement of applications	Potentially significant cost and risk based on availability of expertise for reworking applications
Examples	Domain Name System servers, Lightweight Directory Access Protocol servers, web servers, firewalls, backup and restore, file and print	Business applications running at multiple locations, such as remote offices, bank branches, and retail stores	Rehosting core business applications, such as SAP*, Oracle eBusiness Suite*, and associated databases	An application written to support unique processes for a single business (often in C++, C, Java)

Figure 2. Ease of migration for different application types

The apps you choose to migrate affect the configuration of the hardware platform. For example, if you are moving an older application that does not take advantage of multithreading to the platform, you do not necessarily need a processor configuration with massive numbers of cores. If you have a business objective that involves analytics and big data, your platform may require large amounts of memory and flexible software-defined storage. If the straight-line performance of a specific, single-threaded business application is paramount, processor clock frequencies and accelerator tools are likely to be more important than the total number of available cores. In other words, the nature, capabilities, and demands presented by the software that your business uses defines your hardware requirements, as well as the code and data conversion methodology you choose.

MIGRATION CONSIDERATIONS

When evaluating a U2L migration plan, IT staff members and decision makers should examine the underlying motivations behind the decision. These factors can impact the strategy underlying the migration planning process, influencing migration opportunities, choices, and potential trade-offs.

Most importantly, before proceeding, you need to understand the types of migrations available, as well as the potential deployment scenarios—these can affect every aspect of the migration planning process.

REASONS FOR MIGRATING

When moving from UNIX to Red Hat Enterprise Linux, the following motivating factors are typical:

- Obtain cost reductions in hardware acquisitions, software license and maintenance costs, OS support, system administration, and power usage.
- Replace a server lease that is ending.
- Contend with expanding business requirements while staying within stringent budget constraints.
- Replace discontinued hardware and software.
- Consolidate servers, applications, or datacenter assets.
- Use new technologies, such as virtualization.
- Improve capacity planning and performance.
- Improve security and stability.

In many cases, a combination of these factors will influence the final decision on the OS migration. While no single motivation may be sufficient to warrant the cost, the sum of the business objectives may be enough to justify the migration. In other cases, a single factor, such as substantial cost savings, is important enough to justify the migration. Organizations, however, should also look ahead and consider the value of establishing a flexible, standards-based IT infrastructure to support agile growth and emerging advanced technologies.

COMMON MIGRATION SCENARIOS

In any migration from one OS to another, several migration scenarios must be closely examined in order to create a plan and conduct a successful migration implementation. This section gives a high-level overview of these primary scenarios, progressing from the simplest to the most complex.

UNIX infrastructure application to Red Hat Enterprise Linux infrastructure application

A relatively common scenario—and one of the simplest migration paths—is moving from an external infrastructure application on UNIX to a comparable infrastructure application running on Red Hat Enterprise Linux (see Figure 3). For instance, a customer running Veritas NetBackup* or IBM Spectrum Protect on UNIX may want to continue to do so after migration.

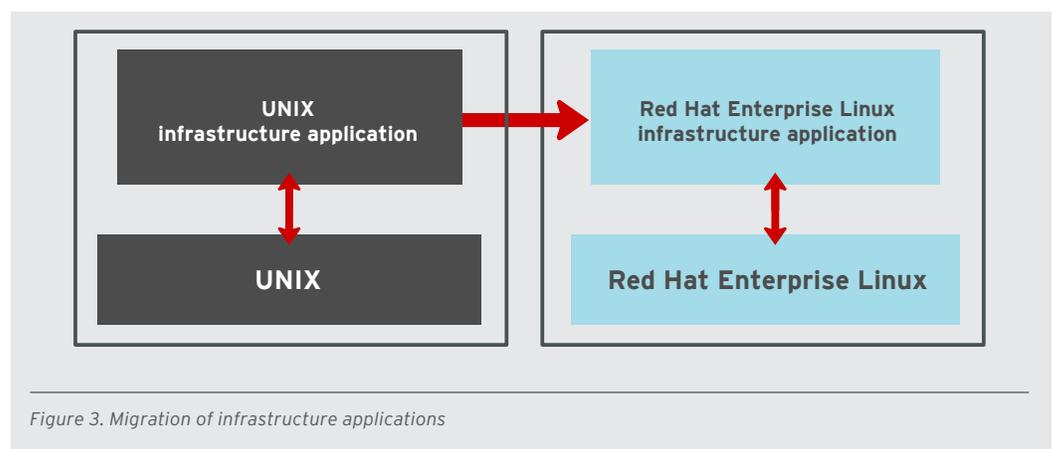
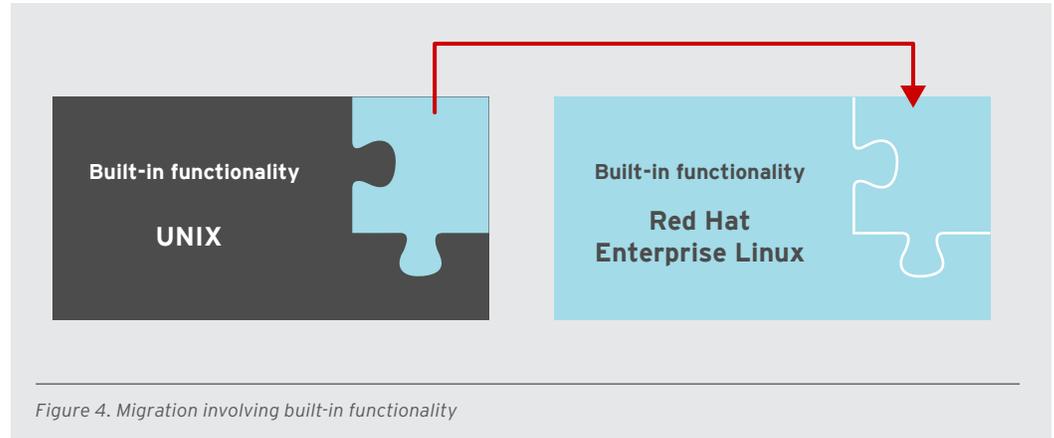


Figure 3. Migration of infrastructure applications

Built-in functionality to built-in functionality

In this scenario, functionality built into UNIX is the same or similar to functions that are built into Red Hat Enterprise Linux (see Figure 4). When functionality is part of both operating systems and works identically (such as Sendmail or NTP), there are few, if any, challenges to migration.



Other infrastructure possibilities

Other infrastructure migration possibilities to consider include:

- **UNIX functionality to Red Hat Enterprise Linux infrastructure application:** In some circumstances, UNIX has built-in functionality that Red Hat Enterprise Linux does not. An additional infrastructure application may be necessary in this scenario to achieve the same functionality in a Red Hat Enterprise Linux environment.
- **UNIX infrastructure application to Red Hat Enterprise Linux built-in functionality:** In this migration scenario, an UNIX infrastructure application is required within a UNIX environment that is not needed with Red Hat Enterprise Linux, which has equivalent, built-in functionality. For example, PowerBroker on AIX for central user management is not needed since Red Hat Enterprise Linux 7 includes identity management capabilities. Substantial cost savings can often be realized because of the wide variety of functionality that is already included in a Red Hat Enterprise Linux subscription.

Office and business-critical application migrations

From the standpoint of migration planning, the kinds of business applications in use within an enterprise fit into three general categories:

- Remote office and retail applications, developed and supported by an independent software vendor (ISV).
- COTS applications that are often available in versions compatible with different operating systems.
- Custom applications that have typically been developed in-house and been in use for many years.

When moving from one core business application on UNIX to the same or similar application on Red Hat Enterprise Linux (see Figure 5), the degree of difficulty depends on the individual circumstances. Options include:

- Identifying an equivalent application that runs on Red Hat Enterprise Linux and migrating the data files to the new environment.
- Porting, converting, or recompiling all or parts of the core business application to run on Red Hat Enterprise Linux.
- Performing the necessary programming to maintain the functionality of the original custom application under Red Hat Enterprise Linux.

As shown earlier in Figure 2, the complexity of the migration rises significantly when moving from basic office applications to COTS and business-critical applications.

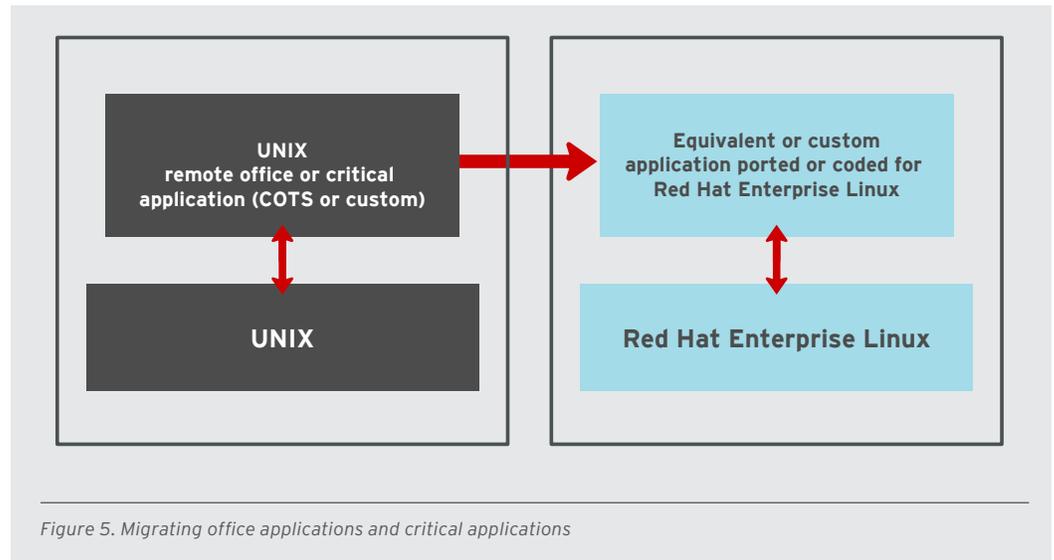


Figure 5. Migrating office applications and critical applications

Custom business-critical applications usually present the most challenging situation unless exceptional care was taken to ensure cross-platform compatibility during their development phase. More detailed information on the best practices to follow when migrating applications to a Linux environment appears in the section “Best practices for a successful migration.”

MIGRATION DEPLOYMENT SCENARIOS

Understanding the possible server-workload deployment scenarios while planning an OS platform migration helps develop the best enterprise architecture for the new environment, resulting in optimal capital cost savings from the migration.

Two of these scenarios, consolidation and dispersion, are discussed in the following sections. Two others, aggregation and cloud migration, are outside of the scope of this document. Red Hat migration services can provide detailed information on these latter two approaches.

CONSOLIDATION

Consolidating underutilized IBM Power* or System p5* systems onto a smaller number of x86-based systems can reduce the datacenter footprint and lower the cost of hardware operation. Customers can take advantage of virtualization to contain individual workloads in virtual machines running Red Hat Enterprise Linux to better utilize compute resources and manage operations. Figure 6 shows an example of a consolidation deployment scenario.

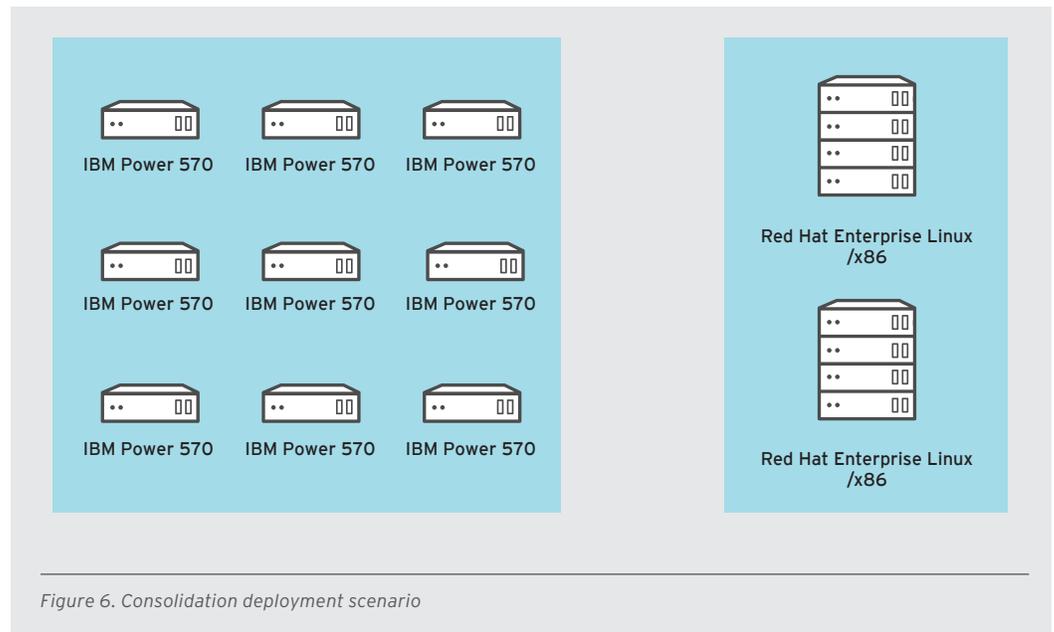


TABLE 1. ADVANTAGES AND DISADVANTAGES OF CONSOLIDATION

ADVANTAGES	DISADVANTAGES
Reduces hardware operational costs	Adds complexity to system administration tasks
Reduces the datacenter footprint	Increases capital costs in scenarios where proprietary virtualization technologies are used
Provides a greater return on investment from the chosen virtualization strategy	Impacts performance to some degree by adding a layer, the hypervisor, between the OS and application
Supports dynamic resource allocation and load balancing	

DISPERSION

This scenario distributes the workloads from one or more Power or System p systems to a smaller number of x86-based systems running Red Hat Enterprise Linux, as shown in Figure 7. Enterprises that are moving more of their IT resources to a Red Hat Enterprise Linux environment often rely on this scenario.

Customers can distribute and scale hardware resources in smaller units across multiple datacenters. While 1U to 4U individual rackmount systems have traditionally been common in this scenario, the use of blades has been growing in recent years. Blade servers provides the customer similar advantages with lower operational costs.

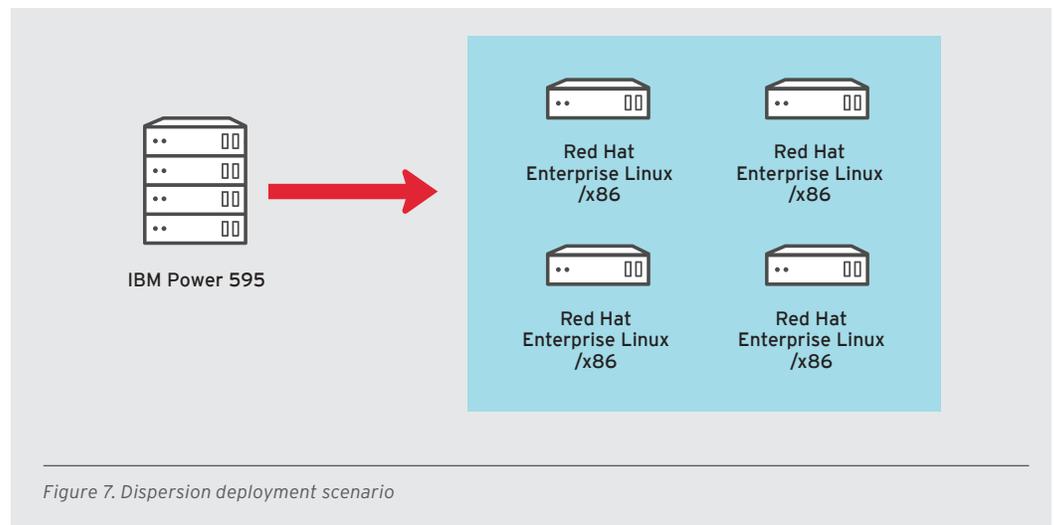


TABLE 2. ADVANTAGES AND DISADVANTAGES OF DISPERSION

ADVANTAGES	DISADVANTAGES
Provides higher performance from the newer x86 hardware technologies	Can result in higher operational costs if not properly planned
Lowers capital cost when scaling hardware resources	Runs the risk of under-utilizing workloads and resources if assets are not optimized properly
Increases flexibility when redeploying resources	

BEST PRACTICES FOR A SUCCESSFUL MIGRATION

This section examines in more detail the stages that take place in typical migration efforts, based on the environment described earlier in Figure 1.

GENERAL GUIDELINES FOR MIGRATING ENVIRONMENTS

Follow these basic guidelines, presented from the simplest to the most complex environments, for handling the platform migration.

Infrastructure applications

Typical examples of applications in this category include DNS, LDAP, web servers, firewalls, backup and restore utilities, and file and print applications.

In many cases, applications in this environment are available in versions that run natively on both architectures, minimizing the need for conversion. Downtime is generally not significant, because switchover can take place immediately. Following an initial migration, the process can be replicated across the enterprise, delivering high value with low cost and risk.

Remote office and retail computing applications

Typical examples of applications in this category include business applications running at multiple locations, such as remote offices, bank branches, and retail stores.

In instances where a business runs the same application in multiple locations, such as remote offices and retail outlets, the cost and risk of migration are distributed across all application migrations. The IT team can effectively complete a single pilot migration, optimize the implementation, and then replicate the migration process across the enterprise servers.

Depending on the solution stack complexity, the cost and risk of this process can vary. If, however, the same software stack is in use across all locations, further migrations should be simple and low risk, providing a high total value for the migration.

Commercial off-the-shelf applications

Typical examples of applications in this category include rehosting core business applications, such as SAP*, Oracle eBusiness Suite*, and associated databases.

The complexity increases for this type of migration because the COTS applications often span multiple layers, including the database layer, application layer, and web services layer. All three layers, in some scenarios, may be hosted in partitions on a large UNIX/RISC server. In other cases, the web or application layers may already be hosted on an Intel processor-based platform. If multiple layers are involved, perform the migration in stages to reduce risk and make it easier to identify and resolve issues that arise during migration.

Code conversion is required if an application does not run natively on Intel architecture. However, shell scripts and customer-coded elements may still need to be converted. Red Hat offers an [application migration toolkit](#) that can help with large-scale application migrations and modernizations.

To begin, perform a detailed assessment of requirements, followed by code conversion. To verify the feasibility of the migration, complete a PoC prior to architecting the solution for scalability and high availability. Size the solution for real-world production workloads and rehearse the migration to optimize the workflow.

Three common scenarios within this environment are:

- **Rehosting an application stack with no software changes:** In this case, only data unique to the enterprise needs to be transferred, simplifying the process.
- **Rehosting an application stack with software version changes:** This includes situations like moving from an earlier version of Oracle to a more current version. Such cases add complexities because of the additional variables involved. Migration tools on the newer version may not be available on the older version. Some coding may be needed to run successfully in the newer version. Consult an experienced migration consultant, if necessary, to ensure a smooth migration process.
- **Rehosting application stacks that require significant upgrades:** This includes cases where changing the software application or underlying database are involved. Software stack upgrades introduce multiple, simultaneous changes that make it difficult to identify issues that arise. To minimize this problem, try to break the migration process down into discrete steps. For example, migrate the existing application tier and then migrate the database.

Business-critical custom applications

Typical examples of applications in this category include an application written to support unique processes for a single business, often written in C++, C, or Java*.

Large-scale code conversion complicates this migration effort. Assessments should consider the availability of source code, porting tools, libraries, documentation, and knowledgeable developers. The time and effort for the code conversion depends largely on the size of the code base.

Automation tools can lessen the degree of difficulty, particularly when porting code from UNIX/RISC environments. Manual coding, however, is typically required, as well as extensive testing to ensure functionality stability, and performance. The migration of custom applications goes more smoothly when the code is modern, well written, and well documented. Older custom applications generally present the most difficulty. They may be poorly documented, heavily patched, and monolithic.

Migration of mainframe applications

The migration of mainframe applications can be relatively simple in many cases. For example, moving a vendor application such as SAP, or a vendor database, such as Oracle Database* or IBM DB2*, is fairly simple. In other situations, particularly when applications running on a mainframe are interdependent, it may be necessary to migrate multiple applications and datasets at the same time, which increases the complexity.

Mainframe migrations to Intel processor-based platforms have been taking place for many years. Tools, expertise, and time-tested processes can minimize the cost, effort, and risk for these migrations. An example of a successful migration of this type by a mainstream organization is discussed in the whitepaper [ReHosting Mainframe Applications on Intel® Xeon® Processor-Based Servers](#).

BUILDING A STANDARD OPERATING ENVIRONMENT

A Standard Operating Environment (SOE) is an organization's standard implementation of the core OS. It can include the base OS, a custom configuration, standard applications used within an organization, software updates, and service packs.

Once an application set has been identified, a standardized build based on an SOE approach is created for rapid and consistent deployment. An SOE build consists of a set of tested hardware, tested software, and configurations deployed on top of Red Hat Enterprise Linux. The SOE build is fully aligned to relevant technical and business requirements, dramatically reducing deployment time, simplifying maintenance, increasing stability, and reducing support and management costs.

PROVISIONING COMPONENTS

If standard hardware is not yet defined, the SOE build can be built on the current corporate standard hardware. Provisioning systems from bare metal to the fully configured SOE build is accomplished through the provisioning feature in Red Hat Satellite.

Provisioning consists of the components listed in Table 3.

TABLE 3. PROVISIONING COMPONENTS

Provisioning configuration	
	Installation methodologies
	Software packages
	Configuration according to security, authentication, storage, and other requirements
Testing	
	Provisioning server setup
	Deployment testing
	Adherence to policy and configuration
Delivery and training	
	Customer's IT staff is trained to deploy and modify the SOE build
	Any remaining customer needs are addressed
	Additional training recommendations are provided
Results	
	SOE builds satisfy customer requirements
Documentation includes:	
	Detailed list of work performed
	Specific procedures, including recommendations for future enhancements or growth
	Links to product-specific manuals
Fully tested provisioning server and provisioning configuration file(s)	
Time-tested and precise methodology frees up resources	

OVERVIEW OF RED HAT SATELLITE

Figure 8 shows the process by which Red Hat Satellite manages SOE builds.

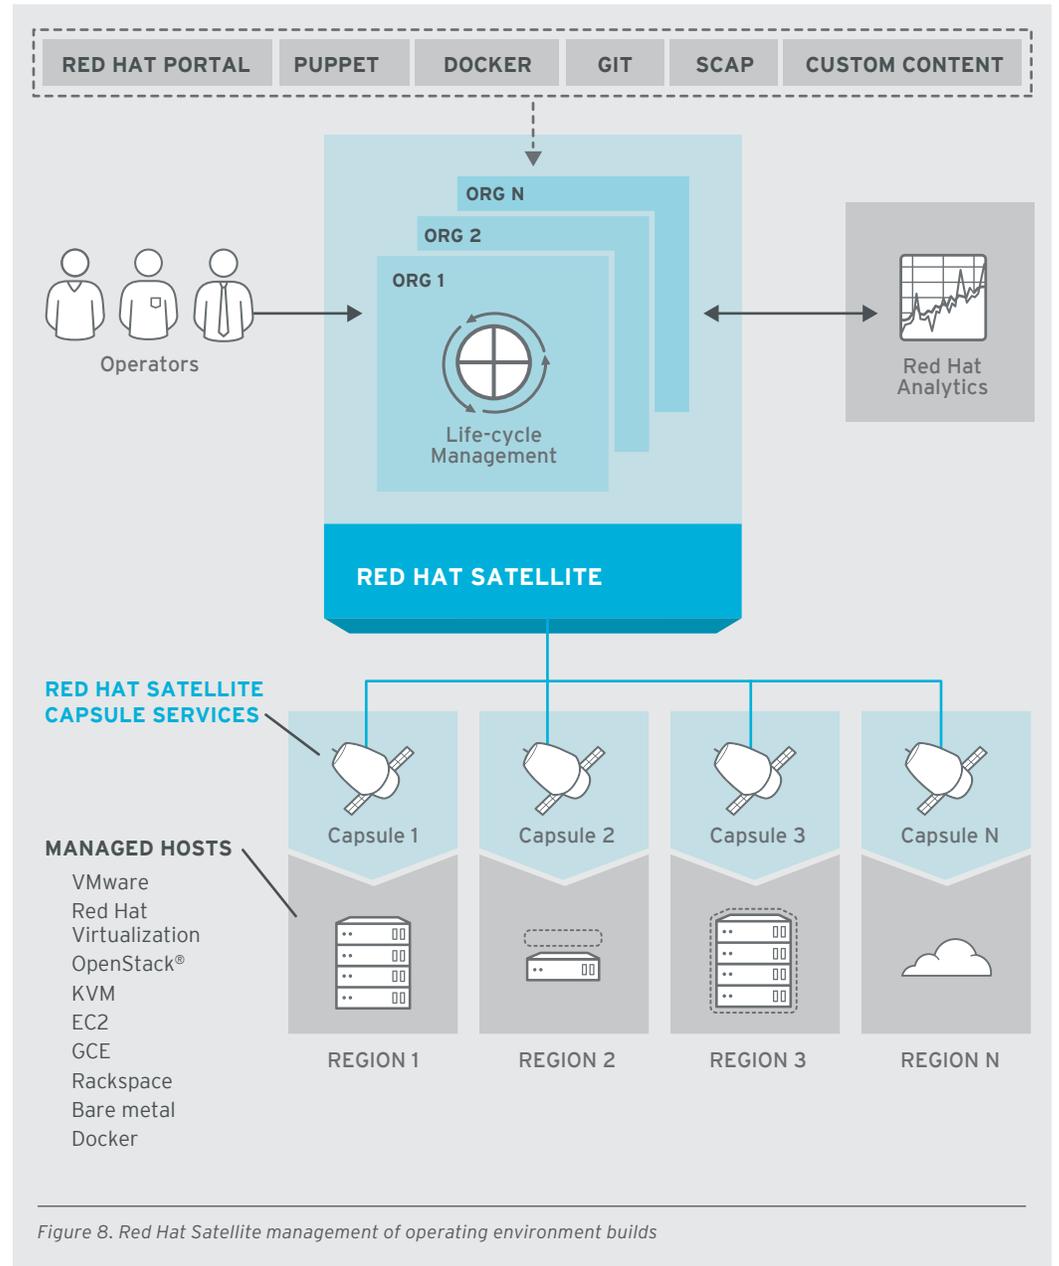


Figure 8. Red Hat Satellite management of operating environment builds

CREATING A STRATEGIC PLANNING ROADMAP

This stage of the strategic migration planning process focuses on bringing together all the information gathered and analyzed so far into a comprehensive strategic planning roadmap.

Creating the strategic planning roadmap includes these steps:

- Compile an inventory of assets.
- Perform a detailed analysis of existing hardware.
- Complete an analysis of the consolidated deployment and virtualization plans.
- Perform a high-level hardware redeployment analysis.
- Create a consolidated risk analysis and risk mitigation plan.
- Develop a training plan.
- Analyze in depth the large, high-complexity applications (optional).
- Estimate the costs in detail.
- Create a master migration roadmap.

CONSOLIDATED ANALYSIS OF EXISTING HARDWARE

The first step in generating the strategic migration roadmap is to perform a detailed analysis of the existing hardware supporting the applications to be moved to the new platform. First, gather the hardware environment data based on an analysis of the functional applications. This includes the following data for development, testing, staging, and deployment environments for each application:

- Number of hosts and processor cores per host.
- Memory requirements.
- Storage and file system requirements.
- Network bandwidth and latency requirements.

The objective is to determine how much processing power, memory, storage, and bandwidth is needed to run the entire set of applications being rehosted. This consolidated view usually represents far more resources than are actually needed to run the set of applications because of the low utilization rates typically present in a datacenter environment.

CONSOLIDATED DEPLOYMENT SCENARIO AND VIRTUALIZATION ANALYSIS

After establishing the aggregate resource needs for the applications targeted for migration, examine application deployment scenarios from the same consolidated perspective. This examination provides a framework for the creation of application groups with common deployment scenario requirements and also offers insight into cost savings opportunities based on virtualization within those groups.

This step provides a consolidated view of how the applications to be migrated map into the new server environment(s).

First, create application deployment groups based upon the preferred application deployment scenarios discussed earlier, keeping in mind the existing application dependencies. This results in one or two application groupings—consolidation and/or dispersion.

Once you establish the grouping data, identify the target hardware profiles for each of the groups. This typically involves working with a set of preferred original equipment manufacturer (OEM) partners, such as IBM, Dell, HP, Cisco, or others. The idea is to create a small number of common system architectures to which the applications targeted for migration can be mapped. View a list of [examples of migration benefits and processes](#).

The fundamental goal is to create as few common system deployment architectures as possible in order to reduce hardware procurement and management costs via standardization. Normally there is at least one system architecture per deployment grouping, but this is not necessarily a requirement. Some organizations have standardized on a single deployment architecture for all migrated applications, regardless of deployment scenario.

Next, perform a high-level virtualization analysis based on the groupings just created. This optional step is highly recommended, depending on a particular organization's policies around virtualization. The virtualization analysis examines several factors for each existing application deployment, including:

- Application SLAs.
- Average and peak hardware utilization rates—processor cores, memory, disk, and bandwidth.
- Physical location of applications—determining which applications are in which datacenter.
- Virtualization limitations—considering ISV support and regulatory and compliance issues.
- Operational type—including development, testing, and production.
- Security and network segmentation—establishing the physical security zone in which each application should reside.
- High availability and disaster recovery requirements.
- Clustering requirements or limitations.
- Specialized hardware requirements—including storage area networks (SANs), InfiniBand*, and other hardware.
- Power and cooling requirements.

These factors determine which applications can and cannot be virtualized—and which virtualized application instances can and cannot be hosted on the same physical machine. The end result of this analysis is a deployment and virtualization map showing a possible arrangement of specific virtual application instances to specific physical machine system architectures.

HIGH-LEVEL HARDWARE REDEPLOYMENT ANALYSIS

As soon as the migration team understands how the migrated applications are likely to be deployed, they can examine possibilities for the redeployment of a subset of the hardware that the applications are currently running. This activity provides an opportunity to offset some of the costs of migration. For example, there may be a set of database instances running on mid-sized IBM Power or System p machines that can be migrated to a Red Hat Enterprise Linux x86 environment. The existing Power or System p machines may then be redeployed into a larger existing database cluster that is not migrated at this time. This process may sound unusual, given that one of the primary cost savings of migrating to Red Hat Enterprise Linux is achieved by eliminating expensive Power and System p boxes, but experience demonstrates that redeployed servers can result in huge cost savings, particularly in situations where the redeployed hardware can no longer be purchased, such as with end-of-life situations, but the hardware is still required to run certain business-critical applications.

This redeployment provides additional capacity without adding new hardware cost and adds savings to the total migration cost estimate.

CONSOLIDATED RISK ANALYSIS AND RISK MITIGATION PLAN UPDATE

In this step, the migration team examines the combined risk factors that were identified in the previous phases of the migration planning process. Additional consideration is provided for any new risks that have been identified in the first three steps of this phase.

This analysis identifies combinations of risks that were previously unknown and could affect the migration. For instance, you may have decided earlier in this process to migrate a large, high-complexity application identified during planning (macro-level difficulty analysis) of the functional application analysis phase. That recommendation may have been based on the risk(s) examined, resulting in a mitigation strategy in the readiness and risk analysis phase that helped determine that the risk was worth it. However, after examining the consolidated deployment scenarios, you may determine that there is additional risk in virtualizing this application. Therefore, the risk migration plan is updated to account for this new risk.

It may also be necessary to update the list of applications targeted for migration based upon these new risk factors and mitigation strategies. This master migration list is used in the detailed cost estimation step.

TRAINING PLAN CREATION

Now that target applications have been identified for migration, optimal physical deployment architectures decided, and the level of organizational readiness understood—from the readiness and risk analysis phase—the next step is to put together a final training plan.

The goal of this step is to identify staff members that need to be trained and the specific training curriculum needed. This almost certainly involves additional Red Hat Enterprise Linux training but may also involve other ISV software training and OEM hardware training from other vendors. Staff members can attend classes that are publicly available on an open enrollment basis or classes can be delivered on-site, depending on specific needs. A set of customized workshops is also available that can be delivered on-site to address topics that are not covered by existing course offerings.

DEEP-DIVE ANALYSIS OF LARGE, HIGH-COMPLEXITY APPLICATIONS (OPTIONAL)

At this point, you may want to go back and revisit the list of large, high-complexity applications that were identified in the macro-level difficulty analysis step of the functional application analysis phase. These applications tend to have the greatest level of uncertainty as to the extent and cost of their migrations. It is often useful to take a closer look at these applications and get a firm grasp on their migration costs before proceeding to the next step—detailed cost estimation. However, this is optional and should be determined on a case-by-case basis.

COST ESTIMATION

Now that all of the information necessary to produce a detailed cost estimate for the entire migration is identified, this step combines the following direct costs and savings in order to come up with a final migration budget estimate:

- Cost of new infrastructure ISV applications necessary to create a Red Hat Enterprise Linux environment comparable to the existing AIX environment.

- Cost of new functional ISV applications necessary to replace existing AIX applications that are not available on Red Hat Enterprise Linux.
- Cost of new Intel architecture-based hardware required to implement each migration deployment architecture.
- Application migration costs.
- Training costs.
- Savings from eliminating proprietary ISV applications and replacing them with open source applications.
- Savings from redeployed hardware.

Keep in mind that this is only an estimate. Actual application migration costs may be more or less. This is not intended to be interpreted as an ROI or TCO analysis because it does not include indirect savings, such as future hardware replacement costs without migrating or operational cost savings.

CREATING A MASTER MIGRATION ROADMAP

In the final step, the master migration roadmap (MMR) is created based upon the input from the previous seven steps. The MMR acts as a project plan that details when, where, and how the migration occurs.

First, analyze and prioritize specific system and application migrations. Relevant factors for prioritizing include capital budget allocation timing, specific business priorities, and datacenter constraints. These factors usually depend on the specific organization, so it is difficult to create a comprehensive list of factors ahead of time.

Once migration priorities have been determined, an actual project timeline is created showing specific dates and durations of the various tasks necessary for a successful migration. This timeline matches specific capital and operational expenditures to quarterly IT budgets, ensuring that migration spend is within budget at all times.

The end result is a migration strategy document drawing on the information captured during the assessment and a project plan with tasks, dates, and expenditures.

SELECTING THE OPTIMAL INTEL® XEON® PROCESSOR-BASED SERVERS

Migration of business-critical enterprise workloads, including large databases, core transactional applications, and business intelligences solutions, can be directed to servers based on Intel® Xeon® Scalable processors, which feature workload optimization capabilities to support the latest hybrid cloud infrastructures and advanced RAS features that maximize server uptime to 99.999% with up to a 1.65 times performance boost, on average.⁴ Tuned for dynamic service delivery, this processor family is well suited to demanding applications, such as analytics, artificial intelligence, HPC, and network transformation.

⁴ "Performance Benchmarks and Configuration Details for Intel Xeon Scalable Processors." Intel. 2017. <https://www.intel.com/content/www/us/en/benchmarks/xeon-scalable-benchmark.html>

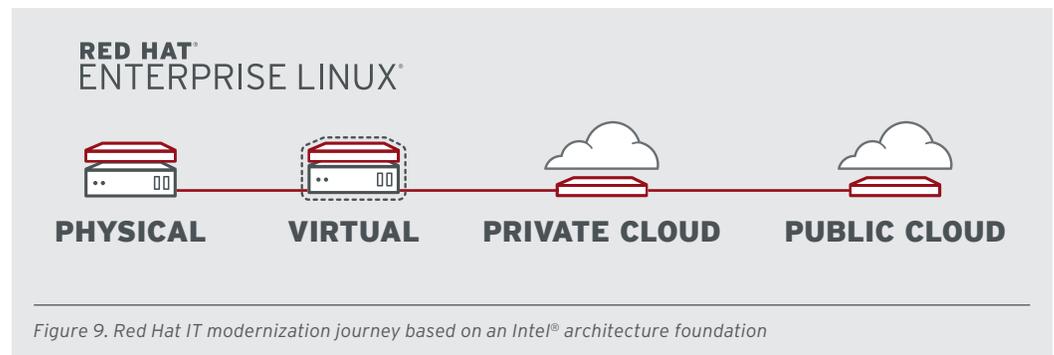
CONTINUING THE RED HAT I.T. MODERNIZATION JOURNEY

Organizations can immediately benefit from building an optimized IT infrastructure in a number of ways:

- Scaling more easily to accommodate system demands.
- Delivering new services more rapidly.
- Creating self-service portals to improve customer experiences.
- Reducing maintenance and administration requirements.
- Cutting data storage expenses.
- Saving on hardware and software expenses.

Once established on a modern, Intel architecture-based hardware platform running Red Hat Enterprise Linux, enterprises can continue on the IT modernization effort by adopting and integrating a wide variety of tools, solutions, and components that will increase agility, elasticity, automation, and scalability. Options include virtualization solutions, software-defined storage, software-defined networks, software-defined infrastructures, management applications, and productivity tools.

As shown in Figure 9, the journey leads from the virtual environment to the benefits of private cloud and public cloud, providing dynamic resource balancing, simplified IT configuration, and many other advantages.



The push toward digital transformation of business enterprises has led to widescale innovation across IT infrastructures, enabling businesses to capitalize on technology advances and the benefits of open standards and open source. Over many years, Intel and Red Hat have worked collaboratively, optimizing Intel architecture platforms running Red Hat Enterprise Linux to perform reliably, securely, and effectively. Reference architectures developed by both companies lead to solid, stable IT infrastructures capable of handling the most demanding enterprise workloads. For example, collaborative work by Dell, Red Hat, and Intel makes it easier for enterprises to deploy the latest software-defined networking and network functions virtualization (NFV) components using the reference architecture provided by the [Intel[®] Open Network Platform](#). The availability of test architectures as a framework is a major factor in cutting development and deployment times for software applications.



TECHNOLOGY DETAIL Best practices and guidelines for optimizing enterprise IT infrastructures

RED HAT MIGRATION SERVICES

Red Hat Consulting offers services to help plan and execute a migration to Red Hat Enterprise Linux. Red Hat consultants are field-tested Linux veterans who guide your staff through migration planning and new deployments of Red Hat Enterprise Linux. Red Hat has decades of experience working with leading companies in financial services, healthcare, telecommunication, and government agencies around the world. Involving a consultant early provides valuable insight into migration best practices, helping you maximize cost savings and efficiencies as the transition progresses, speed time to production, and generate ROI faster.

For more information, visit redhat.com/en/services/consulting.

INTEL RESOURCES AND SERVICES

Intel, along with its partners, can help enterprises migrate from UNIX/RISC to Linux on Intel architecture infrastructures to reduce TCO and increase flexibility and agility to respond to changing business needs. Advanced technologies collaboratively developed by Red Hat and Intel are successfully enabling digital transformation efforts by enterprises seeking the advantages of modern IT infrastructures. Breakthrough developments on software-defined infrastructures expand business opportunities into new realms in which key objectives are efficient resource utilization, demand-driven dynamic reconfiguration, and fluid scalability to accommodate changing workloads.

To learn more about resources available to IT leaders, visit Intel's [New Center of Possibility](#).

ABOUT INTEL

Intel Makes Possible the Most Amazing Experiences of the Future

You may know us for our processors. But we do so much more. Intel invents at the boundaries of technology to make amazing experiences possible for business and society, and for every person on Earth.

Harnessing the capability of the cloud, the ubiquity of the Internet of Things, the latest advances in memory and programmable solutions, and the promise of always-on 5G connectivity, Intel is disrupting industries and solving global challenges. Leading on policy, diversity, inclusion, education, and sustainability, we create value for our stockholders, customers, and society

2200 MISSION COLLEGE BLVD.

SANTA CLARA, CA 95054-1549, USA

PHONE: (408) 765-8080

Intel, the Intel logo, and Xeon are trademarks of Intel Corporation in the U.S. and/or other countries. *Other names and brands may be claimed as the property of others.

ABOUT RED HAT

Red Hat is the world's leading provider of open source software solutions, using a community-powered approach to provide reliable and high-performing cloud, Linux, middleware, storage, and virtualization technologies. Red Hat also offers award-winning support, training, and consulting services. As a connective hub in a global network of enterprises, partners, and open source communities, Red Hat helps create relevant, innovative technologies that liberate resources for growth and prepare customers for the future of IT.

NORTH AMERICA
1 888 REDHAT1

EUROPE, MIDDLE EAST,
AND AFRICA
00800 7334 2835
europe@redhat.com

ASIA PACIFIC
+65 6490 4200
apac@redhat.com

LATIN AMERICA
+54 11 4329 7300
info-latam@redhat.com



facebook.com/redhatinc
[@redhatnews](https://twitter.com/redhatnews)
linkedin.com/company/red-hat

redhat.com
INC0658210_v1_0218

Copyright © 2018 Red Hat, Inc. Red Hat, Red Hat Enterprise Linux, the Shadowman logo, and JBoss are trademarks of Red Hat, Inc., registered in the U.S. and other countries. Linux® is the registered trademark of Linus Torvalds in the U.S. and other countries. The OpenStack word mark and the Square O Design, together or apart, are trademarks or registered trademarks of OpenStack Foundation in the United States and other countries, and are used with the OpenStack Foundation's permission. Red Hat, Inc. is not affiliated with, endorsed by, or sponsored by the OpenStack Foundation or the OpenStack community.