An Oracle White Paper
November 2010

Recommendations for Leveraging the Critical Patch Update and Maintaining a Proper Security Posture

Disclaimer

The following is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described for Oracle's products remains at the sole discretion of Oracle.

Introduction

Software vulnerabilities, flaws in software which might result in compromising the security controls provided by the affected systems, can greatly impact the security posture of an organization. As a result, it is critical that organizations adopt appropriate policies and procedures to deal with security patches to effectively remediate these vulnerabilities when the patches become available. The purpose of this paper is to set forth recommendations to help readers develop appropriate and cost-effective processes for dealing with Oracle security patches (i.e. Critical Patch Updates and Security Alerts) in their organization. These recommendations result from extensive interviews with Oracle customers as well as with various organizations within Oracle who have firsthand knowledge of the Critical Patch Update program in various capacities, including:

1. *Oracle Global Product Security*: The group responsible for the application of the various Oracle Software Security Assurance programs, including overseeing the production of the Critical Patch Update.
2. *Oracle Global Information Technology*: The group responsible for managing and securing Oracle's own IT infrastructure worldwide.
3. *Oracle On Demand* : The outsourcing business unit of Oracle, On Demand is responsible for a broad portfolio of Software-as-a-Service (SaaS) and managed applications enabling its customers to choose how to deploy software based on their own specific needs and budget requirements.

While the purpose of this paper is to provide recommendations that are specific to Oracle's Critical Patch Update, many of these recommendations can be applied to dealing with security patches from other commercial products and vendors. Finally, the mention of third party products in this white paper does not imply a commercial or technical endorsement by Oracle or any of the authors of the white paper. Such mentions are provided for illustrative purposes only.

Finally, note that some of the policies discussed in this document may not be fully applicable to newly acquired product lines (e.g. Sun Microsystems product line). For example, the inclusion of new product lines in the Critical Patch Update program may take some time. In addition, certain variations in policies may be required because of the nature of the technology being brought to Oracle (e.g. policies specific to open source products).

.

"Patch management is no longer optional; however, in the absence of industry best practices, many IT organizations struggle with how much to do and how to do it well."

Ronni J. Colville, Gartner.
Getting Back to Basics on Patch Management,
24 August 2009 ID:G00170521

Summary of Oracle's security vulnerability remediation policies

The very first step an organization needs to take when attempting to define a formal process for dealing with security patches is understanding the practices of its vendors as they relate to the disclosure of the vulnerabilities and the issuance of the fixes. The purpose of this section is to present Oracle's vulnerability remediation programs, namely the Critical Patch Update and the Security Alert programs. The most relevant aspects of these programs, such as the frequency of releases of security fixes, the order in which security bugs are fixed, and the policies governing what information is being disclosed by Oracle will be discussed in detail in this paper.

Critical Patch Updates

In January 2005, the Critical Patch Update (CPU) became Oracle's primary mechanism for the release of security patches for all its products. Today, the CPU program has vocation to provide security fixes for hundreds of different Oracle products. The program is designed to address two strategic goals:

(1) Providing Oracle customers with a cost effective security vulnerability remediation program, and

(2) Maintaining the best possible security posture for Oracle customers before and after the release of the security fixes by Oracle.

Predictability of the Critical Patch Update

A key aspect of Oracle's Critical Patch Update program is its predictability. Oracle releases CPUs on a quarterly basis. The CPU schedule for the next year is posted on the Critical Patch Updates and Security Alerts page on Oracle Technology Network (OTN)[1].

The CPU program was developed with input from Oracle's Security Customer Advisory Council (SCAC). SCAC representatives felt that the frequency with which CPUs are issued (quarterly) and the predictability of the CPU releases (fixed schedule) were necessary to allow Oracle customers to develop

[1] http://www.oracle.com/technetwork/topics/security/alerts-086861.html

a repeatable and cost-effective process for patching their Oracle systems while at the same time, maintaining a proper security posture.

Schedule of the Critical Patch Update

Since the inception of the Critical Patch Update program, Critical Patch Updates were released on the Tuesdays closest to the 15th of the months of January, April, July, October. However, starting in January 2011, the Critical Patch Updates will be released on the Tuesdays closest to the 17th of the months of January, April, July, and October. The Critical Patch Updates and Security Alerts page on Oracle's web site always list the dates of release for the next four Critical Patch Updates, thus effectively providing a one year notice to customers.

On the Thursday before the release of each CPU, a Pre-Release Advisory is published by Oracle. Both the Pre-Release Advisory and the CPU Release Documentation are posted on the Critical Patch Updates and Security Alerts page on Oracle's web site located at http://www.oracle.com/technetwork/topics/security/alerts-086861.html. The content of the CPU Advisory and Pre-Release Notification is discussed later in this paper.

↳ **Useful Tip!**

Oracle customers can choose to subscribe to security notifications from Oracle. Security notifications subscribers will receive an e-mail reminder at the time of the publication of each CPU. Instructions to subscribe to Oracle security notifications can be found on Oracle's web site at http://www.oracle.com/technetwork/topics/security/alerts-086861.html.

Premier Support Customers will also see a message informing them of the availability of the CPU when they log onto the My Oracle Support portal (http://support.oracle.com).

Cumulative nature of the security fixes included in the Critical Patch Update

Critical Patch Updates are cumulative for many Oracle products. This means that, for these products, a CPU includes new security fixes as well as all previously released CPU fixes for this particular platform and version combination. The main benefit of cumulative CPUs is that it allows customers to quickly and easily "catch up" to current security release level by only applying the most recent CPU. A detailed explanation of Oracle's remediation policies, as well as a list of all product families for which CPUs are cumulative can be found on the Security Vulnerability Fixing Policy and Process page[2] on Oracle's web site.

[2] http://www.oracle.com/technetwork/topics/security/alerts-086861.html

Critical Patch Update testing by Oracle

Due to the potential impact of faulty patches, the quality of security fixes is of the outmost importance to Oracle. The objectives of CPU testing in particular are to:
(1) prevent regressions due to application of CPU patches
(2) ensure that the vulnerabilities are addressed effectively (eliminating incorrect or incomplete fixes).

In order to provide enough time for thoroughly testing the database CPU security fixes, the initial selection of the new fixes to be included in the next CPU takes place 15 weeks before this CPU's scheduled release date. During this period, Oracle first creates backports for all individual fixes for the various version-platform combinations currently supported. Fixes are then individually tested and assembled with the other fixes being included in the Critical Patch Update in "early combinations". About 8 weeks before the publication of the next Critical Patch Update, these early combinations are tested to ensure that no regressions have been introduced, and that fixes are fixed correctly and completely.

At the completion of the previous phase, and with about 6 weeks to go before the publication of the Critical Patch Update, Oracle starts testing to ensure that regression issues are not found with other products that use core Application Server and Database products. The "early combinations" are provided to other development organizations to ensure that database CPUs do not create problems across the Oracle stack. For example, throughout this 6 weeks period, the development organizations for E-Business Suite, PeopleSoft, Siebel, On Demand, etc. perform tests with the database CPU early combinations to ensure that database fixes will not create undesirable behaviors with the applications. It is also during this period that Oracle performs additional vulnerability fix testing to ensure that all newly addressed vulnerabilities are effectively addressed in all version-platform combinations of the CPU.

Finally, during the last 4 weeks leading to the release of each CPU, CPU early combinations are provided to a number of organizations within Oracle, including Technical Support to perform installation testing. It is during this period that Oracle ensures that the CPUs install as expected. It is also during this period that the Technical Support organization becomes familiar with the CPU and reviews the CPU documentation for accuracy.

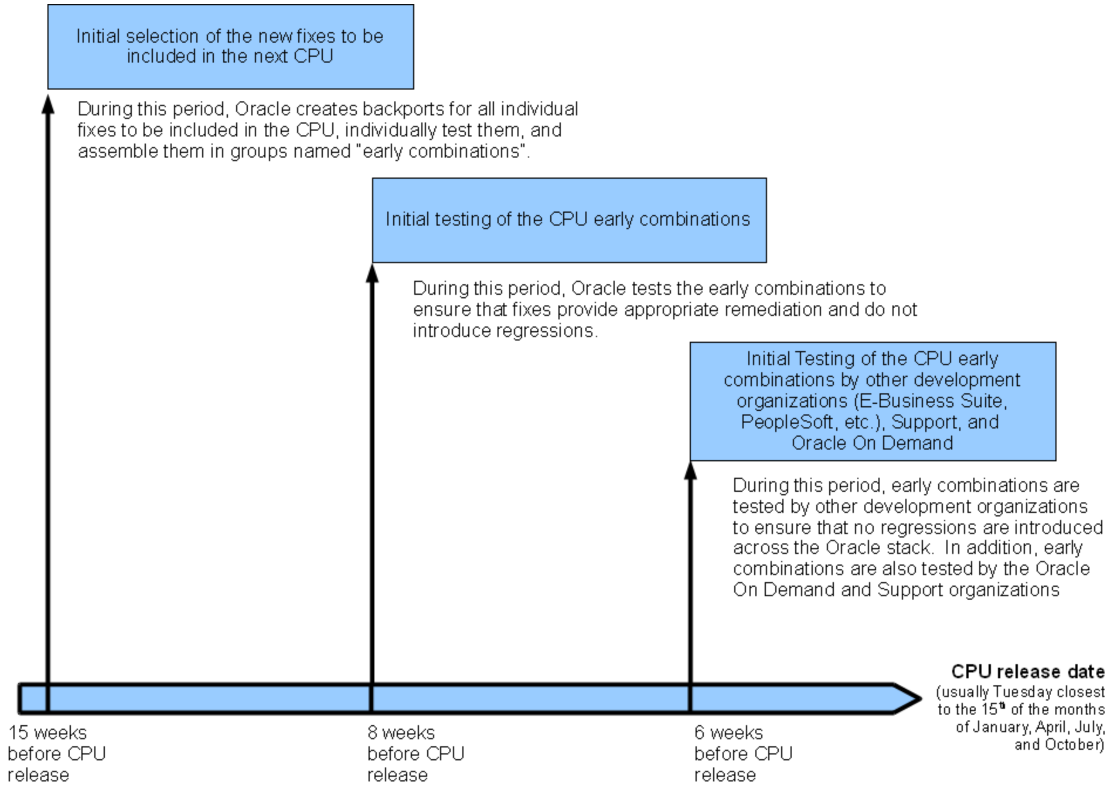
Initial selection of the new fixes to be included in the next CPU
During this period, Oracle creates backports for all individual fixes to be included in the CPU, individually test them, and assemble them in groups named "early combinations".
Initial testing of the CPU early combinations
During this period, Oracle tests the early combinations to ensure that fixes provide appropriate remediation and do not introduce regressions.
Initial Testing of the CPU early combinations by other development organizations (E-Business Suite, PeopleSoft, etc.), Support, and Oracle On Demand
During this period, early combinations are tested by other development organizations to ensure that no regressions are introduced across the Oracle stack. In addition, early combinations are also tested by the Oracle On Demand and Support organizations
CPU release date
(usually Tuesday closest to the 15th of the months of January, April, July, and October)
15 weeks before CPU release
8 weeks before CPU release
6 weeks before CPU release

Figure 1: Typical testing timeline for CPU patches

Security Alerts

In certain circumstances, the existence of particularly critical vulnerabilities or the publication of exploit code “in the wild” may require Oracle to release out of cycle security fixes. In such an instance, Oracle can exceptionally issue off-schedule Security Alerts (outside of the regularly scheduled CPUs), which would include security fixes or workaround instructions to address particularly significant threats to Oracle customers.

As of November 2010, and since the introduction of the Critical Patch Update program in January 2005, Oracle used the Security Alert mechanism in only three occasions (for more information, see the Security Alerts for CVE-2008-3257, CVE-2010-0073, and CVE-2010-0886[3]).

[3] http://www.oracle.com/technology/deploy/security/alerts.htm

↳ **Useful Tip!**

Subscribing to security notifications from Oracle will ensure that you are notified in case of release of a Security Alert. Instructions to subscribe to Oracle security notifications can be found on Oracle's web site at http://www.oracle.com/technetwork/topics/security/securityemail-090378.html

Oracle Security Vulnerability Disclosure policies

Maintaining the security posture of *all* customers is Oracle's foremost priority with its security disclosure policies. It has been Oracle's experience that the acknowledgement of a possible security bug will result in more undesirable scrutiny by malicious hackers. Furthermore, the company has observed that malicious attackers will also often attempt to analyze security patches to determine the nature of the fixed vulnerabilities in an attempt to develop exploit code. In recent years, Oracle and the whole security industry have observed the acceleration of the development of malicious code as a result of the publication of security patches and even anti-malware signatures. As a result, the company has adopted the following disclosure principles.

1 *Do no harm*! Oracle does not want to disclose any information that an attacker might use to develop a successful exploit against an Oracle product. Many of the policies below are logical implications of this principle.

2 Oracle does not comment on published reports of alleged Oracle vulnerabilities.

3 Oracle will not announce security fixes until they are available for all affected and supported product version and platform combinations.

4 Oracle will not provide workaround instructions unless the workaround provide for the complete elimination of the vulnerability on all supported platforms and versions combinations.

5 Oracle will not provide additional information about the specifics of vulnerabilities beyond what is provided in the Critical Patch Update Advisory, the Security Alert or documents referenced from these two documents.

6 Oracle provides all customers with the same information, at the same time, in order to protect all customers equally. Oracle will not provide advance notification or "insider information" on CPU or Security Alerts to individual customers.

7 Oracle does not develop or distribute active exploit code (or "proof of concept code") for vulnerabilities in its products.

↳ **Useful Tip!**

Third party information about patched vulnerabilities (including statements made on security blogs and web sites) may be erroneous. Oracle's security patch documentation (such as the CPU Advisory) should be the only authoritative source of information on Oracle security vulnerabilities. Oracle customers can contact Oracle Support if they have specific questions about information reported in the CPU Documentation, including the risk matrices.

Furthermore, external reports of alleged Oracle vulnerabilities are often incorrect

> in many important ways and the purported workarounds or mitigations instructions in these reports have often been found to be ineffective and/or known to cause serious regressions.

How are security vulnerabilities discovered?

Security defects can be discovered through internal means (such as automated testing or developer peer review) or reported by external parties (such as customers, technology partners, or security researchers). Oracle generally fixes security vulnerabilities in severity order regardless of how the vulnerabilities are discovered. This means that the vulnerabilities that create the greatest risk to customers are fixed first. Of course, the likelihood that a vulnerability is publicly known impacts its fix prioritization because public knowledge of a vulnerability generally does increase the risk of exploit code being created and released "in the wild."

In recent years, the increasing use of automated tools by Oracle has had an impact on the proportion of security defects that are discovered internally versus those reported by external sources. This is because the increasing use of automated tools allows Oracle to find proportionally more security defects internally. Of all the security defects found in 2009 (as of September 15th 2009), 87% were found internally, 10% were reported by customers, and 3% were found externally by non-customers. This last group consists of defects reported to Oracle by security researchers, Oracle-specific defects posted directly to the Internet, and problems in third party products/code that are included with Oracle products. For more information about this topic see Darius Wiles' blog entry[4] dated October 6, 2009.

Once a vulnerability has been discovered or reported to Oracle, it is analyzed to assess -- for example -- what components or interfaces are affected, which components of the code are vulnerable, how the vulnerability can be exploited, how it can be eliminated, what other mitigation measures are possible, etc. Throughout this process, Oracle may engage with the reporter of the vulnerability in instances when the vulnerability was externally discovered. It is the only instance when Oracle shares vulnerability-related information with a third party. However, the scope of these communications is limited to the vulnerability that was reported by the third party.

At the end of this initial analysis phase, a fix for the vulnerability is produced. This fix is then reviewed by developers and security experts before being included in the main code line (that is if the vulnerability is not limited to previous versions of the affected component) where full regression tests are run. Once the inclusion in the main code line is complete, the fix is ported to all future patch sets. The fix may then be scheduled for inclusion in a future Critical Patch Update following the process detailed in section 1.1.4.

[4] http://blogs.oracle.com/security/2009/10/security_defect_testing.html

Useful Tip! Oracle customers or Oracle partners can use My Oracle Support to submit a Service Request on any potential security vulnerability affecting an Oracle product. Alternatively, an email can be sent to secalert_us@oracle.com; however, when using e-mail, Oracle encourages the use of email encryption (Oracle's encryption key can be found at http://www.oracle.com/technetwork/topics/security/encryptionkey-090208.html).

Key requirements for organizations for developing repeatable and cost-effective patching policies

In the preparation of this white paper, interviews have found that the most common technical and environmental factors affecting an organization's approach to security patching were:

- Size and complexity of the environment: systematic patching of all systems is typically not possible in large environments because of costs and resources required. The more diverse the environment, the more difficult it is to patch all systems because of the significant testing effort and requirement for platform-specific knowledge.
- Absence of buy-in or excessive objections to patching because of production needs: organizations who have not obtained the buy-in from the various stakeholders find it very difficult to justify the production impact of security patching efforts (e.g. downtime)
- Existence of relevant security policies or controls: organizations with a good understanding of the security controls available around sensitive systems typically can make better security patching decisions (for example, they can delay the application of security patches because of the existence of external mitigation measures). Furthermore, the size of the user pool, the number of active accounts on systems and how these systems are accessed have a great impact in the priority given to security patching.

Existence of technical inventory and configuration controls

The amount of control exercised by an organization over its IT environment significantly impacts the organization's ability to effectively deal with security patching in a cost-effective manner. There are two aspects to this control that can positively impact an organization with its security patching effort:

(1) The ability of the organization to gather accurate systems inventory information, and

(2) The level of control the organization has over systems configurations deployed in its environment.

This is because an organization needs to have a good understanding of (and control over) its systems inventory in order to effectively deal with security patching. The greater the number of different configurations an organization needs to deal with, the more difficult it is for the organization to maintain all these disparate systems and keep them at current patch level.

Inventory information should not be limited to software and hardware version information. Nor should it be limited to a single or a few elements of the technology stack (e.g. configuration information limited to database or operating system versions). It is important that the organization collects enough relevant information to allow for a complete and accurate replication of the production environment for testing and validation purposes. Another important requirement is for the organization to map business processes and functions to production systems so that the operational impact of altering or patching IT systems can be easily and accurately understood.

Furthermore, the consistent use of organization-approved configurations (e.g. "configuration baselines") can significantly ease patching effort. By limiting the heterogeneity of the systems that potentially need to be patched; the organization will require less effort for developing test plans. The documentation of the patching procedures, as well as the effort required for information sharing between the technical staff will also be greatly reduced. Finally, the more homogeneous the environment, the simpler it is to perform an accurate risk assessment prior to making patching decisions. In that context, a more homogeneous technical environment is also easier to secure.

Note that organizations should not rely on policies alone to enforce baseline or corporate-approved configurations in their production environment. Periodic scans of the environment are required to identify deviations to "official" baselines (and systems inventory). Such scans are required to detect unauthorized installation of software or services (which can result in increasing the attack surface of the organization), introduction of new systems (such as introduction of rogue wireless access points which create additional risks for the organization), or undue configuration changes made by the IT staff without proper authorization, including the creation of unneeded accounts.

In summary, the more information the organization has about its environment, the more control it has over deployed systems configurations, the more effective the organization can be with patching and risk management.

Understanding of organizational tolerance for risks

The risk exposure an organization is willing to assume, as well as the nature of the threats the organization faces, must be weighed against production requirements when security patching decisions are required.

It is critical that all the relevant parties, including the Lines of Business (LOBs) owners, who may be impacted by security activities, be involved in the definition (and quantification) of the risks an organization is willing to assume. The involvement of the various stakeholders in the definition phase of the security policies is required in order for them to understand -in principle- the production impacts resulting from normal security activities, such as patching. Interviews in the preparation of this white paper have found that the most common reason for "stonewalling" the application of otherwise required security patches resulted from the lack of initial involvement of LOBs whose systems were to be impacted temporarily by the application of the patches. However, the process for assessing technical risks posed by vulnerabilities is better left handled by a specialized security group because of the unique knowledge and skill set required to understand the impact of security flaws.

From an organizational perspective, three basic principles need to be followed so as to create a favorable ground for developing an effective and repeatable patching process:

1. Patching policies need to be documented. They must refer to the appropriate organizational security and audit policies, and all relevant stakeholders must be involved in their definition.
2. The various stakeholders need to develop an accurate understanding of what it takes to test and apply patches as well as have a general understanding of the possible impact resulting from non-application of the security patches.
3. Clear accountability must be maintained any time a deviation to policy is requested, required, denied, or obtained. The standard audit rules should always apply, including accountability, transparency, and evidence principles.

Ability to leverage existing security policies and controls

Poor security controls or lax enforcement of generally accepted security principles can have significant impact on the priority that an organization will need to give to security patching in order to maintain a reasonable security posture. For example, the more people there are with access (or active accounts) to systems via public (i.e. insecure) and direct TCP/IP connection, the greater the priority that should be given to security patching because of the exposure of these systems. For example, databases with large number of users (or connected applications) with read access privileges need to be patched particularly quickly.

In many instances, when database access is limited to a single application server and a handful of properly vetted IT personnel, organizations may choose to only apply patch sets (when they are available) over Critical Patch Updates. The judicious use of network controls (through firewalls and network routers) can help organization reduce their risks by reducing the exposure of sensitive systems. However, the use of such network access controls cannot mitigate all threats, and organizations need to enforce proper monitoring of their systems and consider security patching when appropriate.

Generally, the extent to which external mitigation controls are implemented in the organization should determine the aggressiveness with which Critical Patch Updates are applied, for example:

- As a rule of thumb, access to production databases via TCP/IP should be limited to applications and the few IT administrative personnel in charge of the maintenance of the related systems on an exception basis. To be clear, this means that firewalls or routers should block access from all but these accounts so that unauthenticated network attacks would be possible only from these few IT administrators and applications.
- Privileged or developer accounts should be disabled in production unless a specific action is required. When the account is activated, preferably by the information security group, the activities should be tracked and the account disabled once the assigned task has been completed. Access to production systems by developers and non IT personnel should be avoided as much as possible.

- In addition, as it relates to protecting database server, organizations should try -as much as possible- to avoid managing user accounts on the database side. Access control policies should be enforced through business applications (and additional database security controls should be used to avoid "application bypass" issues) and the deployment of development tools and frameworks to certain environments (production, Citrix servers, etc.) should be restricted.
- In addition, files and directories should be protected with well defined access rights in an attempt to reflect "need to know" principle and prevent "nice to know" curiosity.
- Organizational "hardening" effort, resulting in the definition of reasonably secure base configurations, should not be limited to the database, but extended to the operating systems, applications, and network layers.
- It is also critical that organizations periodically audit their systems to control what accounts exist on these systems, and for example, remove improperly created administrative accounts and eliminate inactive accounts.

IT organizations should maintain "situational awareness" that is a comprehensive understanding of their IT risk. The existence of mitigation measures can determine whether vulnerabilities fixed in the Critical Patch Updates can be exploited by a malicious user. Organizations with more sophisticated risk assessment and countermeasures tend to have well-defined patching policies (including backup and recovery plans). Organizations with good situational awareness may decide to postpone fixing vulnerabilities because their successful exploitation will be prevented by means other than patching. However, such organizations need to recognize that any change in outside controls may result in allowing the exploitation of the vulnerabilities. That is to say, they are giving up maintaining a "Security in Depth" posture. Therefore proper documentation is recommended to keep track of any such operational risks. The organizations then need to regularly review these operational risks and continuously apply general audit principles. Ultimately, it is expected that the increasing burden of documenting and keeping track of a number of vulnerabilities that have not been patched because of the availability of external controls against ever changing infrastructure changes, will cause the organizations to apply patches so as to "reset" its risk posture (and regain its security in depth posture) against the queue of known but unpatched vulnerabilities.

Useful Tip!

A number of security guides and checklists are available from Oracle. These documents can be used to help organizations develop secure baselines for their Oracle deployments. The following table (Table 1) lists a few of these resources.

TABLE 1 PARTIAL LIST OF THE SECURITY GUIDES AND CHECKLISTS AVAILABLE TO ORACLE CUSTOMERS

Oracle Database Server	
Oracle Database Security Guide 11g Release 2 (11.2) (OTN)	http://download.oracle.com/docs/cd/E11882_01/network.112/e10574/toc.htm
Oracle Database Security Guide 11g Release 1 (11.1) (OTN)	http://download.oracle.com/docs/cd/B28359_01/network.111/b28531/toc.htm
Oracle Database Security Guide 10g Release 2 (10.2) (OTN)	http://download.oracle.com/docs/cd/B19306_01/network.102/b14266/toc.htm
Oracle Database Security Guide 10g Release 1 (10.1) (OTN)	http://download.oracle.com/docs/cd/B14117_01/network.101/b10773/toc.htm
Oracle Database Security Checklist (OTN)	http://www.oracle.com/technology/deploy/security/pdf/twp_security_checklist_db_database.pdf
Oracle Fusion Middleware	
Securing a Production Environment for Oracle WebLogic Server 11g Release 1 (10.3.1)	http://download.oracle.com/docs/cd/E12839_01/web.1111/e13705/toc.htm
Configuring A Secure Oracle9i Application Server Environment, a lockdown guide (OTN)	http://www.oracle.com/technology/deploy/security/oracle9iAS/pdf/securingias.pdf
Oracle Applica ion Server Security Guide 10g Release 2 (10.1.2) (OTN)	http://download.oracle.com/docs/cd/B14099_09/core.1012/b13999/toc.htm
Oracle Applica ion Server 10g Security Guide 10g (9.0.4) (OTN)	http://download.oracle.com/docs/cd/B10464_05/core.904/b10377/toc.htm
Oracle E-Business Suite	
Best Practices for Securing Oracle E-Business Suite Release 12 white paper, February 2007 (MetaLink)	https://support.oracle.com/CSP/main/ar icle?cmd=show&type=NOT&id=403537.1
Security Forum on Oracle.com	
My Oracle Forum on Security Best Practices	http://myforums.oracle.com/jive3/thread.jspa? hreadID=46514

Ability to make educated patching decisions

While the systematic application of Critical Patch Updates against the entire production and test environment is highly desirable and recommended by Oracle, the reality of production requirements, the pressure to meet service level requirements, and the cost of repeated wide-scale patching may prevent organizations from applying security patches systematically to all affected systems.

Interviews have found a wide range of security patching practices: from systematic application of security patches (because of limited control over the environment and absence of corporate baseline) to justifiable application of patches (resulting from a risk assessment analysis performed at the time of each CPU release). Organizations with control over their environment and a clear understanding of the costs associated with patching can adopt formal security patching policies and procedures that reflect their production requirements and the risk posture that they are willing to assume. These organizations are also well-equipped to make informed decision for skipping patches altogether without significant negative impact of their risk posture. Figure 2 includes a simple decision tree for helping organizations manage patching decisions.

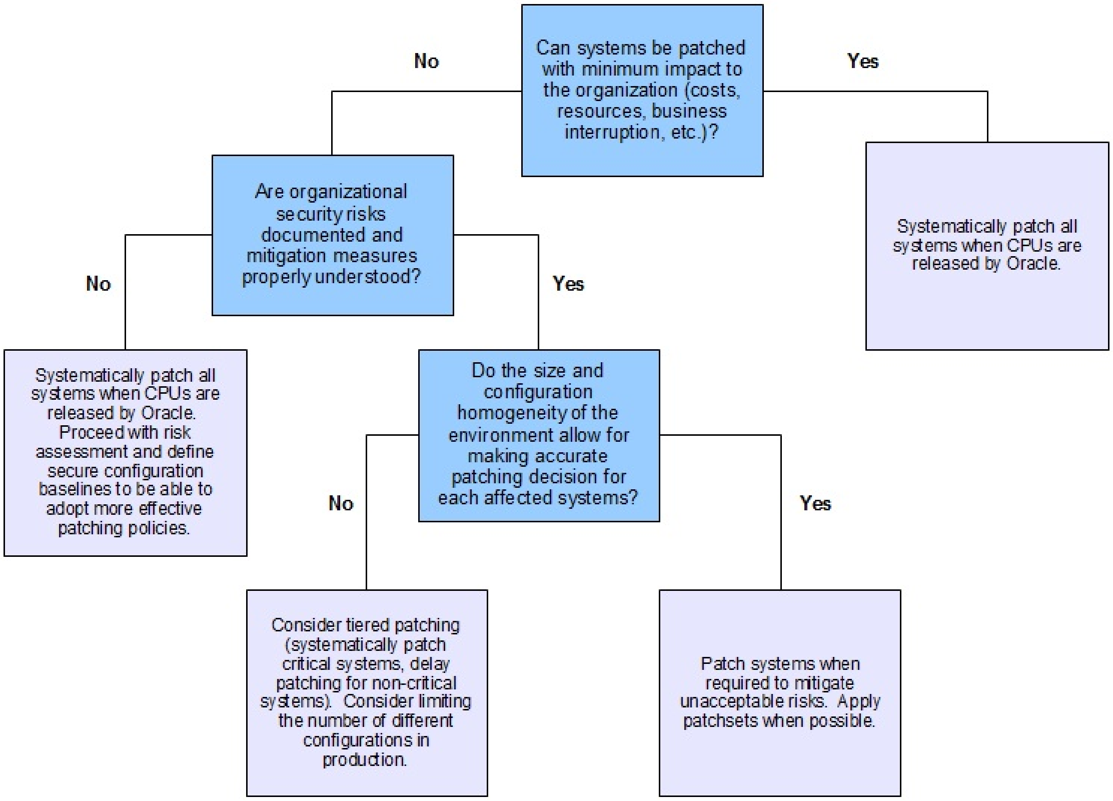
Can systems be patched with minimum impact to the organization (costs, resources, business interruption, etc.)?
No
Yes
Systematically patch all systems when CPUs are released by Oracle.
Are organizational security risks documented and mitigation measures properly understood?
No
Yes
Systematically patch all systems when CPUs are released by Oracle. Proceed with risk assessment and define secure configuration baselines to be able to adopt more effective patching policies.
Do the size and configuration homogeneity of the environment allow for making accurate patching decision for each affected systems?
No
Yes
Consider tiered patching (systematically patch critical systems, delay patching for non-critical systems). Consider limiting the number of different configurations in production.
Patch systems when required to mitigate unacceptable risks. Apply patchsets when possible.

Figure 2: Patching Decision Tree

Furthermore, most interviewed organizations stressed the importance of assessing the "real cost" of security patching when making patching decisions. There are 4 components to this cost:

1 Engineering of the new release including QA (i.e. planning phase),

2 QA effort to guarantee functional and non-functional transparency for patched solutions (most expensive part),

3 Reviewing application and business deployment plans (and adjusting them if needed to prioritize CPU application across the targeted systems), and

4 Actual CPU deployment cost including the cost of business interruption and/or service interruption.

User activities at the time of the publication of the Pre-Release Notification by Oracle

Oracle issues a CPU Pre-release Notification on the Thursday before the publication of each Critical Patch Update (on the following Tuesday). The publication of the Pre-Release Notification by Oracle should serve as a trigger to initiate the organization's patching procedures.

Content of Oracle's CPU pre-release notification

The goal of the CPU Pre-Release Notification is to remind customers about the upcoming CPU and provide them with a high level introduction into the nature of the upcoming security fixes.

The Pre-Release Notification provides a summary of the risk matrices in the upcoming CPU Advisory and contains the following information:

- Product and Component names and versions affected by new fixes
- Number of fixes for each product family
- Maximum CVSS score for each product family
- Number of remotely exploitable vulnerabilities fixed for each product family
- Number of client-only fixes (if any).

Note that, because the testing of the CPU doesn't end with the publication of the pre-release advisory (testing continues until very near to the actual publication of the CPU), the content of the actual CPU may be slightly different from the information provided in the pre-release notification.

In rare instances, a small number of fixes which were announced in the pre-release notification may be dropped from the actual CPU because of concerns discovered during late phase testing in one or more platform and version combinations. In such instances, the announcement of the affected vulnerabilities will be delayed until a future CPU when fixes are published for all platform and version combinations.

Pre-approvals & buy-ins

The process for seeking all the necessary organizational approvals for the application of the CPU in production environment should be documented and repeatable. The publication by Oracle of the Pre-Release Notification can mark the beginning of the organization's approval process. This is because the Pre-Release Notification can be used -at a high level- to determine whether patching will be required in the organization, and the scope of the patching effort (list of all the systems that are potential candidates for the upcoming CPU).

Planning activities following the publication of the Critical Patch Update by Oracle

The CPU Advisory is posted on the Critical Patch Updates and Security Alerts page[5] on Oracle's web site on the CPU release day at about 1PM Pacific Time. This document should serve as the cornerstone of the organization's patching effort as it lists the new fixes introduced by the CPU, their severity, and where to get the appropriate fixes and application instructions.

Interpretation of the CPU Advisory & its risk matrices

Once the CPU Advisory is published, organizations should consult the Advisory to determine whether new fixes apply to their systems, and assess the severity of the vulnerabilities for which these fixes are released. Note again that the CPU Advisory is the only authoritative source of vulnerability information about Oracle products as Oracle doesn't provide "additional information" to any third party organizations.

Content of the CPU Advisory

The CPU Advisory is a detailed document that includes the following information:

- *Supported Products and Components Affected*: This section lists the products and the affected versions for which new security fixes are released in the Critical Patch Update. In addition, the Advisory also lists the product area for the listed products. Customers should refer to the risk matrix under each corresponding product area to get more information about the vulnerabilities affecting them.
- *Patch Availability Table and Risk Matrices*: This section lists the product groups affected by the CPU and indicates whether the patches are cumulative. It also provides a link to the appropriate risk matrices (published in the CPU Advisory itself) and installation documentation (typically available as Support Notes and therefore accessible through the My Oracle Support system).
- *Workarounds instructions*: If available, a section containing workaround instructions for specific security flaws will be published in the CPU advisory. These workaround instructions can help an organization mitigate the risks associated with a specific flaw until a patch is made available by Oracle. Note that Oracle does not provide workaround instructions unless they are fully effective workarounds for all customers on all supported platform and version combinations.
- *Unsupported Products & De-Supported Versions*: This section is intended to provide a warning to customers that unsupported and de-supported products do not typically receive security patches, and therefore may be vulnerable to attacks. Oracle recommends that customers remain on supported versions so as to continue to get security patches and yield the benefits of Oracle's ongoing security

[5] http://www.oracle.com/technetwork/topics/security/alerts-086861.html

assurance activities. Note however that Critical Patch Update patches are available to customers who have purchased Extended Support under the Lifetime Support Policy. For more information, see the Oracle Lifetime Support Policy available on Oracle's web site[6]. Oracle's Technical Support Policies are also available on Oracle's web site at http://www.oracle.com/us/support/assurance/fixing-policies/index.html

- *Credit Statement*: This section lists the individuals (and their organizations) who have submitted security vulnerability reports to Oracle and adhered to responsible disclosure practices by working with Oracle until a fix was completed to address the vulnerabilities they discovered.
- *Modification History*: This section lists all the changes made to the CPU Advisory document since its initial publication (on the CPU release date).

Use of the risk matrices to assess risks in the environment

Oracle has designed the risk matrices to allow customers to determine the risks addressed by the fixes distributed in association with these documents. In a risk matrix, each row describes a single vulnerability that has been addressed by fixes in the distribution (Critical Patch Update or Security Alert).

Row entries can essentially be divided into (1) product information, (2) attack vector information and (3) CVSS information. This information should be adequate to allow customers to determine the degree of risk to their organization for the associated vulnerability. The columns of the risk matrices should be interpreted as follows:

Product Information

- *Risk Matrix Title*: This defines the product or product suite for this matrix and establishes a context for products described in this risk matrix. Some examples are Oracle Database Risk Matrix, Oracle Application Server Risk Matrix or Oracle Enterprise Manager Risk Matrix.
- *CVE#*: This is the industry standard identifier of the vulnerability and is provided by the Common Vulnerability and Exposures group at http://cve.mitre.org/ Note that CVE, the National Vulnerability Database, and many other sites, often include more information than is provided in the Oracle Advisories. Oracle recommends that all security vulnerability information be obtained from Oracle (see section 1.3).
- *Component*: This is the high level component affected by the vulnerability. Examples for the Database might be Oracle Spatial or Application Express. One should consider this component at risk if it is enabled for use even if the component is not used. Thus, for example, even if a

[6] http://www.oracle.com/us/support/lifetime-support/index.html

customer is not using Oracle Spatial, they should consider themselves at risk if Spatial is enabled but not used by the customer applications.

Attack Vector Information

- *Protocol*: This is the protocol over which the vulnerability can be exploited. Reported protocols typically include TCP/IP such as HTTP or Oracle Net. If the attack is launched via the Operating System then the reported protocol is designated "Local" or "Local Login". In some instances, it is possible to mitigate the vulnerability on the affected systems by blocking or limiting connections using the reported protocol.
- *Package and/or Privilege Required*: This is either a subcomponent under the component or the privilege required to launch an attack. When this column contains a privilege, the nature of the privilege required will often be very important in determining risk. For example, if the privilege required is "Session only", meaning that only a logon is required, the risk is much greater than if the privilege was reported as "create table".
- *Last affected patch set (per supported release):* This column may be one of the most difficult columns to understand in the CPU or Security Alert Advisory because it first requires the reader to know which product releases are currently supported. This is determined by reviewing the Lifetime Support policy for the product, which defines the supported versions of a product, and by reviewing the Software Error Correction Support policy for the specific product, which details how software fixes are delivered (see for example Support Note 209768.1 for Oracle Database Server). Once the reader has established which releases are supported, there are two possibilities:
    - For every supported release version that does NOT have a specified patchset listed in the column, there is no vulnerability.
    - For every release version that does have a specified patchset listed in the column, the specified patchset and all prior patchsets are vulnerable.

  For example, let's assume that the reader has determined that Oracle Database version 10.2 is supported per Lifetime Support Policy. If there are no entries in the form 10.2.x.x.x specified in the "Last Affected Patch Set" column, then no version of Database 10.2 has the vulnerability in question. If patchsets 10.2.0.3 and patchsets 10.2.0.4 are supported but the column only reported "10.2.0.3", then this would mean that Database version 10.2.0.4 is not vulnerable (but 10.2.0.3 is).
- *Remote Exploit without Authentication*: This column is derived from the Access Vector and Authentication CVSS columns (see section 4.1.2.3 "CVSS Information")

Except for the notes (last column of the risk matrix), the remainder of the columns contain CVSS information.

CVSS Information

CVSS (Common Vulnerability Scoring System) is a standard of the FIRST organization. CVSS information has two categories. The first pertains to the ease of access by an attacker while the second pertains to the impact of the vulnerability.

Access values include the following:

- *Access vector*: The values reported by Oracle are "Network", which means an attack can occur over the network, and "Local", which means that only local attacks are possible (i.e. attacker has physical access to the machine). Generally, local only attacks may be considered lower risk in instanced where the IT staff is trusted (and has been properly vetted).
- *Access complexity*: This column reports on the difficulty of launching an attack that has already been created. It does not pertain to the degree of difficulty in creating the attack. Values are Low, Medium and High. Low means easy access and typically requires no or low levels of privilege. Medium means higher privileges required and may require "social engineering" e.g. cross site scripting. High means high privileges or special situations or difficult timing (such as very narrow attack windows). Note that historically, the values reported by Oracle typically were "Medium" and "High" depending on the degree of social engineering (e.g. Cross Site Scripting is usually Medium) or the level of privilege needed (e.g. Database CREATE TABLE privilege would be medium, ability to perform BACKUP would be High).
- *Authentication*: This column indicates whether authentication is required in order to exploit the vulnerability. Possible values are : "None", "Single Authentication" or "Multiple Authentications". In practice, Oracle rarely has vulnerabilities that require Multiple Authentications.

Impact values in CVSS area include:

- Confidentiality: Unauthorized disclosure of data
- Integrity: Unauthorized create/update/delete of data
- Availability: Unauthorized denial of service

For each of the type of impact (confidentiality, integrity and availability), the following values may be reported:

- None: No impact
- Partial: Impact on some of the data of the product
- Partial+: the exploit affects a wide range of resources, e.g. all database tables, or compromises an entire application or subsystem. The addition of the Partial+ rating does *not* change the CVSS base metric scoring system.
- Complete: Impact on all of the data of the box where the product executes.

More information on CVSS can be found on the CVSS web site located at http://www.first.org/cvss/. However, note that:

1 Oracle makes a strict application of the CVSS standard when computing the Base Score. CVSS base scores are computed via a formula published at the FIRST site where the CVSS standard is defined. For the purpose of the computation of the CVSS Base Score, Oracle considers Partial+ and Partial to be the same. Some organizations may choose to deviate from the CVSS standard and discretely inflate the Base Score when a value of Partial + is reported by Oracle.

2 The possibility of a complete takeover of the box where the affected product executes (down to the OS layer) will exist when a value of “Complete” is reported in the Confidentiality, Integrity and Availability columns for a given vulnerability in the risk matrix. The takeover of a product (wide compromise limited to the affected application) is possible when a value of “Partial+” is reported in all three columns.

↳ Useful Tip!

A basic understanding of the CVSS Base Score brings you a long way in understanding the severity of the vulnerabilities fixed in a given CPU or Security Alert. For example:

- A CVSS Base Score greater than 7.5 indicates a compromise of the box where the product executes.
- A value of 10.0 indicates an easy, over the network, unauthenticated and full takeover of the box where the product executes.
- A value of 9.0 typically indicates a relatively easy, over the network, and full takeover of the box where the product executes, but a low privilege authentication is required.
- A value of 7.5 with no “Complete” value reported in the impact columns of the risk matrix indicates an easy, over the network unauthenticated takeover of the product. (A value of 6.5 would be reported if all impacts were reported as “Partial+” and if authentication was required.)

Assessing the “real” risk for the organization

While Oracle is trying to provide as much information as possible to allow organizations to assess the criticality of the vulnerabilities fixed in the Critical Patch Update or Security Alert, the estimation of the actual risk created by these vulnerabilities require a good understanding of the operational environment of the organization. The actual risk for a particular site depends on a number of factors. For example, an internal only attack whose impact is only Availability (Denial of Service) might not be considered high risk at many sites given that employees are unlikely to mount an attack that would have little value, would be easy to detect and would likely involve a good chance of discovery and certain prosecution. Also, the impact values are very important. Partial, Partial, Partial impacts may only affect a small part of a product where Partial+, Partial+, Partial+ indicates a takeover of the product. Finally, while the CVSS Base Score does not differentiate between accessibility from the Internet versus Intranet, this may have a significant impact on the risk for a particular organization. This is because an interface only accessible to Intranet users creates probably far less of a risk for the organization than one accessible to the Internet, especially if the product in question is both isolated from the Internet as well as being isolated from all but a few IT personnel in an organization.

Identification of the systems that need to be patched & definition of path test plan

In previous sections, we have stressed the need for organizations to maintain accurate inventory information, and if possible to promote the use of common corporate-approved configuration baselines in order to limit the heterogeneity of systems found in production. The use of third party systems management products such as CA Unicenter, IBM Tivoli, Oracle Enterprise Manager and others may help organizations with maintaining proper inventory. In addition, Premier Support Customers who have enabled the Collector (Oracle Configuration Manager) can see on the My Oracle Support portal an accurate inventory of their systems (for which the Collector was enabled) and receive patch recommendations and other technical advice to keep their systems running smoothly[7].

Ideally, organizations will have a firm control over the configurations of their most critical systems, and will be able to replicate a similar configuration in a test environment to test the application of the Critical Patch Updates before their deployment in production. For organizations dealing with a large number of heterogeneous systems, a phased approach may be required for testing and deploying CPUs. The priority given to patching individual systems will need to be determined based on the criticality of each system.

The final step in this planning phase is to determine whether applying the patch is required for each system, and the timeframe for the application of the CPU. Organizations interviewed in the preparation of this white paper reported that common options included:

- Apply CPU quickly: e.g. the criticality of the security bugs and the importance of the systems warrant immediate application of the CPU.
- Apply CPU deferred: e.g. the criticality of the security bugs, mitigated by external controls, and the importance of the systems warrant application of the CPU in the next scheduled maintenance window.
- Apply Patchset: e.g. the low criticality of the security bugs, and the relative importance of the systems, along with the soon release of an Oracle patchset, which include functionality fixes that are desirable for the organization, warrants delaying the application of the CPU in favor of applying the upcoming patchset.
- Apply new release: e.g. the low criticality of the security bugs that are completely mitigated by external security controls, along with the expectation of a new release by Oracle in the coming months, which include significant functional enhancements that are highly desirable for the organization, warrant delaying the application of the CPU in favor of upgrading the systems to the new release.

[7] For more information, see the Collector tab on the My Oracle Support portal.

CPU vs. Patchset decision

In the previous section, we stated that organizations may be faced at times with decision to apply patchsets vs. CPUs. This is because security fixes are included in patchsets (fixes are included in main code line and patch sets before inclusion into CPUs) for non-terminal releases[8].

In such an instance the voluntary "skipping" of a CPU to apply an upcoming patchset is acceptable provided that the additional time of operation with unpatched vulnerabilities doesn't induce excessive risks for the organization. Again, the availability of well understood security controls around Oracle systems may allow organizations to effectively mitigate these vulnerabilities long enough to wait for the application of the upcoming patchset.

Patch Set Updates (PSUs) vs. traditional CPU patches

Patch Set Updates (PSUs) are an enhanced patch offering, which were introduced with the July 2009 CPU for Database versions 10.2.0.4 and later. PSUs vary from traditional CPU patches in that they include security fixes and other recommended bug fixes (non-security fixes).

PSUs include low risk/high value fixes, such as fixes for critical technical issues (wrong results, corruptions, hangs, etc.), and fixes issues that have been encountered by large number of customers. PSUs are available only for certain Oracle products and product versions. PSUs are also not available on the Windows platform, but the PSU content is included in the Windows Database bundle patches.

The application of PSUs results in the introduction of a new baseline version identified by the 5th place version number (e.g. Oracle Database Server 10.2.0.4.2). Organizations need to make a determination as to which patching mechanism they will commit to (PSU vs. traditional CPU format), because while PSUs are released under the normal CPU schedule and contain the same security content as traditional CPU patches, PSUs employ a different patching mechanism. Once a PSU has been installed, the recommended way to get future security content is to apply subsequent PSUs. Reverting from PSU back to CPU, while possible, would require significant effort, and is therefore not recommended by Oracle.

Support Note 854428.1[9] provides a detailed explanation of the PSUs along with a list of Oracle products for which PSUs are produced.

[8] More information about Oracle's policies for fixing vulnerabilities is posted on http://www.oracle.com/us/support/assurance/fixing-policies/index.html

[9] https://support.oracle.com/CSP/main/article?cmd=show&id=854428.1&type=NOT

Validation of backup procedures

Prior to the actual application of the CPU, organizations should ensure the effectiveness of their backup systems and procedures. In other words, organizations need to backup their systems and ensure that these backups are valid, and can be restored without loss of data.

Backup media should be kept available until such a time that patched systems have been put back into production and accepted by their respective systems owner.

Testing and deployment activities by users

In this section, we will briefly discuss the activities that need to take place in the organization for the testing and deployment of CPUs.

Testing

Validation of testing requirements

The testing of CPUs in test environments before the actual application of the CPUs in production should be designed to ensure that the performance of the systems are not negatively impacted by the application of the CPUs and that their application doesn't result in breaking "applications." While database CPUs do not typically include functional changes, changes of database behavior may impact higher level applications. CPUs should be applied in test environments to provide the application group enough time to ensure that they do not result in negatively impacting the operation of the applications. In addition, the application of CPUs in test environments provides the organization with the ability to test CPU roll back procedures to test backout procedures and validate the time it would take to revert back the patch in production if it became necessary.

Definition of testing environment

Test environments should be a close representation of production environments in order to provide for accurate testing. Active test systems are also necessary for testing the application (and in the absence of active test systems, a well-defined test plan should aim at testing all major functions of the application).

Definition of the tests

The existence of a test plan that touches the different aspects of the production systems is extremely important. Test plans should include a standard test for database, application and connectivity. Full regression testing based on allowable risk and complexity of the environment may be required for the most critical systems in the organization. A good understanding of Oracle's testing practices is required to draft an effective plan so as to focus on the areas that have not been previously tested by Oracle (see section 1.1.4 of this white paper).

Documentation of the results

Organizations should keep all log files created during patch application, and also maintain all log files of checks that were accomplished during testing. Such documentation is important for keeping track of which tests passed without issues, and possibly debug issues in production if encountered.

Deployment of security patches

Considerations when "skipping" CPUs

Organizations with large heterogeneous environments may face unique challenges resulting from the testing and deployment cycles of the CPUs, for example, patching for the previous CPU may still be ongoing in production as the new CPU is being released by Oracle. In certain instances, an accurate understanding of the organizational risks and controls available to the organization might just make it more practical to apply every other CPU. In the instances where CPUs are cumulative (that is they include all fixes released in previous CPUs), organizations can quickly "catch up", even when they have skipped past CPUs, by applying the most recent CPU.

Considerations when previously current on CPU

An enhanced CPU patch format was introduced with the July 2007 CPU for customers on Oracle Database Server 10.2.0.3 and later on Unix. The n-apply CPU provides a number of benefits, including the ability to skip previously installed patches, which saves on install downtime and minimizes change to the system. Support Note 438314.1[10] "*Critical Patch Update - Introduction to Database n-Apply CPUs*" provides detailed information on how to apply n-apply patches.

Considerations when working with previously cloned systems

A number of organizations have reported that they accelerated CPU deployments by working on previously cloned environments. Using this approach, new Oracle homes are created and cloned as a copy of the existing homes. CPUs can then be applied to the new homes first before swithing them with the existing homes. This approach allows for applying the patches while the database is up, and then during the maintenance window, the outage time will be limited to just switching the Oracle homes. The initial home is also a good backup of the binaries to be used if rollback is required. This approach is obviously an option for organizations when hardware and storage are readily available. The benefit of this approach is that the application of the CPU will cause minimum interruption (when switching to the previously cloned and newly patched system) and rollback involves minimum risk (by returning to the unpatched system that was previously in production).

[10] https://support.oracle.com/CSP/main/article?cmd=show&id=438314.1&type=NOT

Considerations when working in RAC environment

Database CPU patches are rolling RAC installable, enabling the application of the patch without requiring the database to be taken down. Customers should refer to My Oracle Support Note 244241.1 "Rolling Patch - OPatch Support for RAC" for more information.

Finding out how to apply the security patches

The processes for applying patches vary per product families. Customers should always refer to the appropriate documentation as a starting point before engaging in the deployment of the patches. The appropriate documentation is always referenced in the Critical Patch Update or Security Alert Advisories. The following chart provides examples that can be used to locate the proper documentation.

TABLE 2 DOCUMENTATION AVAILABLE WITH CRITICAL PATCH UPDATES AND SECURITY ALERTS

PRODUCT GROUP	AVAILABLE DOCUMENTATION	COMMENTS
Oracle Database Server	Patch Availability Document referenced in each CPU Advisory.	In addition, the following documents may be consulted: • Introduction to Database n-Apply CPUs [ID 438314.1] • Patch Set Updates for Oracle Products [ID 854428.1]
Oracle Application Server and Fusion Middleware 11g	Oracle Application Server 10g Examples for Critical Patch Updates - Plus FMW 11g [ID 405972.1]	
Oracle Collaboration Suite	Applying Critical Patch Updates to Collaboration Suite 10g [ID 559534.1]	
Oracle E-Business Suite	Knowledge management document listed in each Critical patch Update	For example, see "Oracle E-Business Suite Releases 11i and 12 Critical Patch Update Knowledge Document (April 2010)" [ID 985896.1]
PeopleSoft Enterprise **JDEdwards Enterprise**	Knowledge management document listed in each Critical patch Update	For example, see "Oracle Cri ical Patch Update Advisory - April 2010" [ID 1077118.1]
Industry Applications Product Suites (e.g. Life Sciences, Retail, Communications, etc.)	Knowledge management document listed in each Critical patch Update	For example, see "Critical Patch Update April 2010 for Oracle Communications Products" [ID 1076933.1]; "Oracle Life Sciences Applications Critical Patch Update Note April 2010" [ID 1063379.1]; "Critical Patch Update April 2010 Patch Availability Document for Oracle Retail Products" [ID 1078890.1]

Ensuring availability of patches in the next CPU (On Request model for Application Server and Database Server)

Once patches have been deployed and systems have been accepted back into production, the last step for the organization is to ensure the continued availability of security patches from Oracle with the next Critical Patch Update. For the Application Server and Database Server product families, Oracle systematically creates patches only for those platform and version combinations, which based on historical data, customers are likely to download[11].

For historically inactive platform and version combinations of the Oracle Database and Oracle Application Server (that is platform and version combinations for which there were few downloads by customers of the previously released CPUs), Oracle will create patches only if specifically requested by customers. Additional details regarding the products, versions and platforms that will be supported for the next Critical Patch Update and the process for requesting On Request patches are available in the Critical Patch Update Patch Availability Document for Oracle Products.

Customers should refer to the CPU Patch Availability Document for Application Server and Database Server to make sure that Oracle plans to release fixes in the next CPU for their platform and version combinations. If not, these customers should contact Oracle Support in order to get a CPU produced for their products.

[11] See My Oracle Support Note 1060989.1 (http://support.oracle.com/CSP/main/article?cmd=show&type=NOT&id=1060989.1) for more information.

Oracle is committed to developing practices and products that help protect the environment

Recommendations for leveraging the Critical Patch Update

August 2010
Author: Bruce Lowenthal, Eric Maurice
Contributing Authors: Michelle Malcher (IOUG), Lois Price, Darius Wiles

Oracle Corporation
World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065
U.S.A.

Worldwide Inquiries:
Phone: +1.650 506.7000
Fax: +1.650.506.7200
oracle com

SOFTWARE. HARDWARE. COMPLETE.