



How Well Do You Know Documentum? Lessons Learned, Some Obvious, Some Maybe Not So Obvious

Doug Desaulniers

Eisai

NE Documentum User Group Meeting

October 27, 2009





Agenda

- Introductions
- Technical Lessons Learned
- The Important Lessons Learned
- Q&A
- Prediction





Introductions





Eisai

- **Our Mission – Human Health Care (hhc)**
 - *hhc* means we regard patients and their families as the most important participants in the health care process.

- **Our Goal**
 - Our goal is clear: We look for and produce innovative solutions - specialty products - that can satisfy unmet medical needs and contribute to the health and well being of people worldwide. We can achieve this goal independently or in collaboration with other companies.





Eisai

■ Some of Our Products

- **ARICEPT® (*donepezil hydrochloride*)**
 - For the treatment of Alzheimer's disease
- **ACIPHEX® (*rabeprazole sodium*)**
 - For the healing of erosive gastroesophageal reflux disease (GERD)
- **BANZEL® (*rufinamide*)**
 - For adjunctive treatment of seizures associated with Lennox-Gastaut syndrome (LGS) in children 4 years and older and adults
- **ALOXI® (*palonosetron hydrochloride*)**
 - For the prevention of acute and delayed nausea and vomiting associated with initial and repeat courses of chemotherapy
- **DACOGEN® (*decitabine*)**
 - A therapy approved to treat myelodysplastic syndromes (MDS), a group of diseases that result from ineffective blood cell production





Eisai

- **The Company**
 - **Eisai, Inc is the US affiliate of Eisai Company, Ltd.**
 - **Founded in Japan in 1941**
 - **A publicly traded company on the Tokyo Stock Exchange and Osaka Securities Exchange**
 - **Listed in the United States under an American Depositary Receipt with the ticker symbol of ESALY**
 - **FY2008 revenues of \$7.8B**
 - **More than 10,000 employees worldwide**





Doug Desaulniers

- Director, Document Management Systems
 - Joined Eisai in May 2007
- Prior Experience
 - Content Management Strategist for a leading management consulting firm
 - Sr. Manager, Document Management Center of Excellence, Schering-Plough
 - Director, Project Delivery Services, First Consulting Group
- Contact Information
 - doug_desaulniers@eisai.com, 201-627-2341





EDNA & Eddie

- Our Document Management Systems
 - EDNA
 - Supports manufacturing operations
 - Repository for procedures, forms, work instructions, and batch records
 - Hosted in the US
 - Approximately 400 domestic users
 - Eddie
 - Supports clinical development and regulatory submissions
 - Hosted in Japan
 - Approximately 1,400 global users
 - Both are built on DCTM 5.3 and DCM 5.3
 - Eddie uses RenderPerfect for renditions
 - EDNA uses PDF Aqua for renditions





Document Management Systems Department

- Our Philosophy
 - If it were not for the users, we wouldn't be needed
 - A big THANK YOU to the business
- Our Goal
 - Make sure the systems are delivering the intended efficiency to the business processes, i.e., we are getting the anticipated return on our investment
 - Overall System Performance
 - Design (cabinets/folders, security, etc.)





This Presentation

- The Audience Representation
 - Business?
 - IT?
 - Vendors?
 - Who uses DCM?
- The Invitation from CYA
- How Well Do You Know Documentum?
 - And some other things
- Thank You EMC





Technical Lessons Learned





Newer Is Not Always Better





Newer Is Not Always Better

- In addition to Documentum and DCM, there are “supporting” software components in the architecture, such as IE, Java, and Adobe.
- EMC has built Documentum to be highly flexible so it will work with most current versions of the supporting software.
 - A. True
 - B. False





Newer Is Not Always Better

- **Background**
 - Some users need to import multiple documents, or export multiple documents (sequentially, rather than in bulk)
 - The first import or first export worked fine
 - On subsequent attempts, the system just hangs
 - The fix is to shut-down and re-start IE
- **The Cause**
 - Someone knew the version of java being used was “outdated” and installed a newer version on the user’s desktop





Newer Is Not Always Better

- The Lesson (should be obvious)
 - Follow EMC's guidance and make sure the supported versions of the supporting software is installed.





Darned If You Do, Darned If You Don't





Darned If You Do, Darned If You Don't

- After logging on to the system, the user was presented with the following screen.





Darned If You Do, Darned If You Don't

The screenshot shows a Windows Internet Explorer browser window titled "Documentum Compliance Manager - Windows Internet Explorer". The address bar displays the URL: `http://esidmpwp01/dcm/webcomponent/library/searchresultslist/search...`. The page content includes a search bar with the text "recovery" and a list of search results. A JavaScript error message is displayed in the main content area: `appintgevents.js: AppIntgProcessEvent did not handle aEvent(event=ShowBusyCursor)`. The error message is repeated twice. The browser interface includes a menu bar (File, Edit, View, Favorites, Tools, Help), a search bar with the Google logo, and a sidebar with a file explorer view showing folders like "Home Cabinet", "Inbox", and "Subscriptions". The status bar at the bottom shows "Internet" and "100%" zoom.





Darned If You Do, Darned If You Don't

- What was the cause of this?
 - A. Incompatible Java version
 - B. Using IE version 7
 - C. Using IE version 8
 - C. Incompatible Adobe version





Darned If You Do, Darned If You Don't

- **Background**
 - When Documentum starts, the javascript function isAppIntgEnvironment() is executed to determine if the appropriate javascript object model is being used by IE
 - Most IE javascript DOMs return “undefined” when an undefined property is encountered.
 - In this case, the IE javascript DOM was returning “unknown”
- **The Fix**
 - Add “unknown” to the javascript function isAppIntgEnvironment() in the wdk\include\appintgevents.js file





Darned If You Do, Darned If You Don't

- The Lesson (maybe not so obvious)
 - Following EMC's guidance and installing the supported versions of the supporting software does not mean there won't be issues to be researched and resolved.





Use The Tools They Give You





Use The Tools They Give You

- Where is the TCP/IP Window Size controlled?
 - Content Server, Web Server and Client
- What is the default value for the TCP/IP Window Size parameter on the Content Server and Web Server?
 - Depends on the network speed
 - Below 1 Mbps, default is 8KB
 - Between 1-100 Mbps, default is 17KB
 - Greater than 100 Mbps, default is 64KB






Use The Tools They Give You

- **Background**
 - Eddie was upgraded in January 2007 from 4i to 5.3
 - The 4i system was installed in a Unix environment
 - The 5.3 system environment is installed in a Windows OS 2003 environment
 - The most common complaint from US and UK users after the upgrade was “It’s so slow”





Viewing a PDF Document – in US Before & After Upgrade





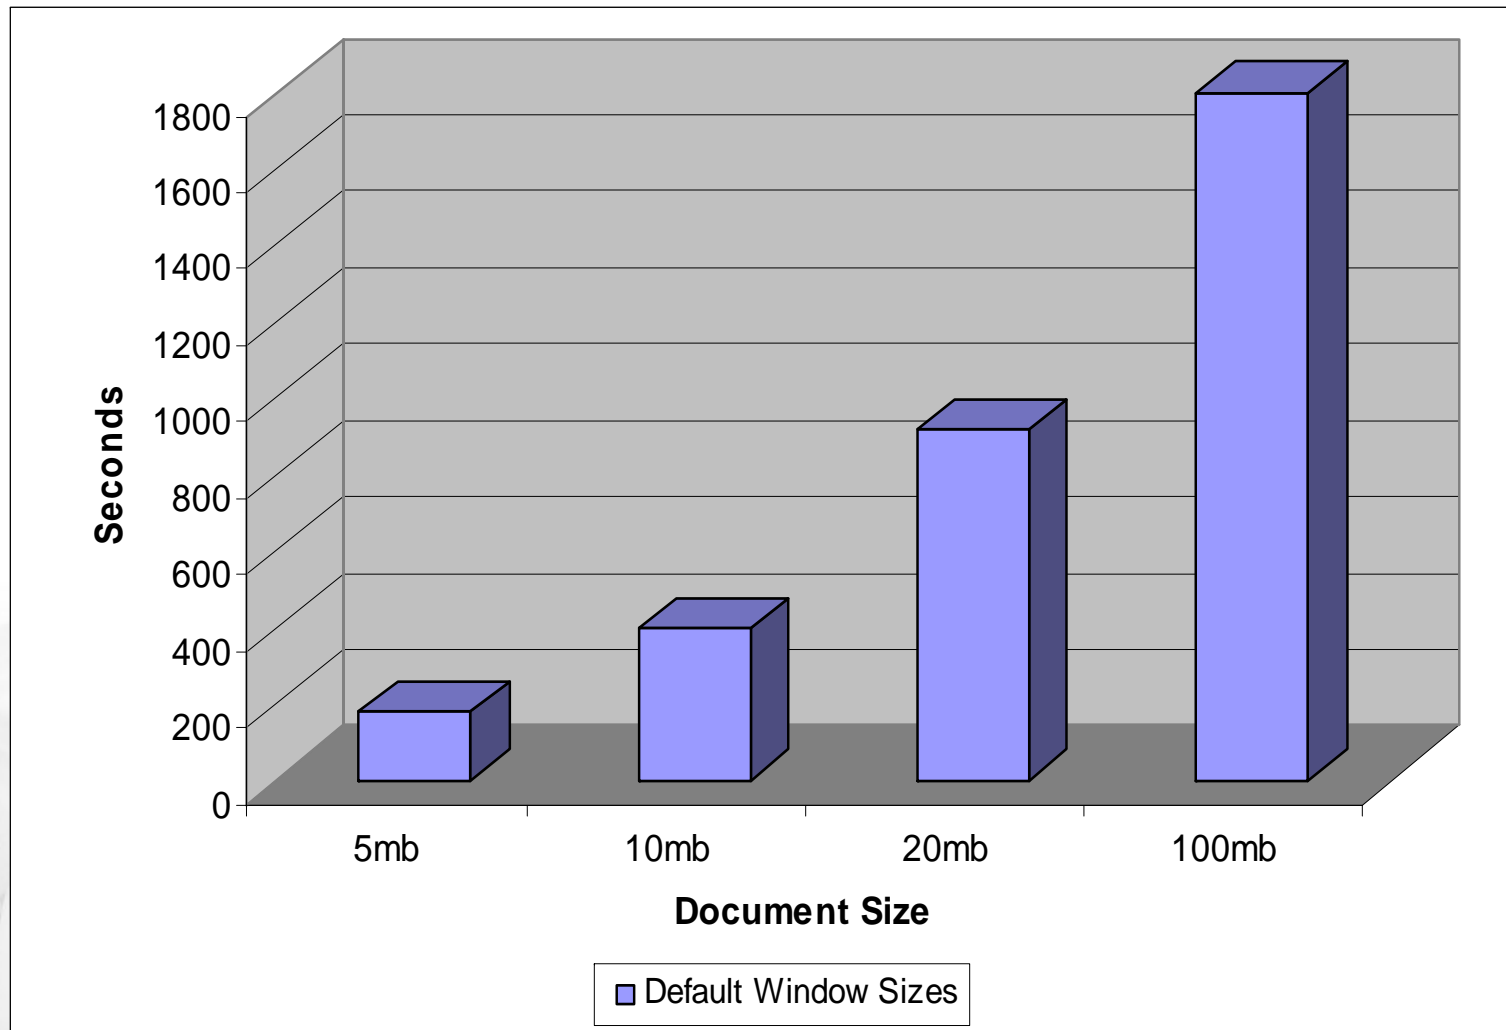
Use The Tools They Give You

- **Background**
 - During the upgrade, the TCP/IP Window Size was set in the Windows registry
 - Rather than using Maximum Segment Size
 - Default registry value was used
 - 8KB






Viewing a PDF Document – Default Window Size



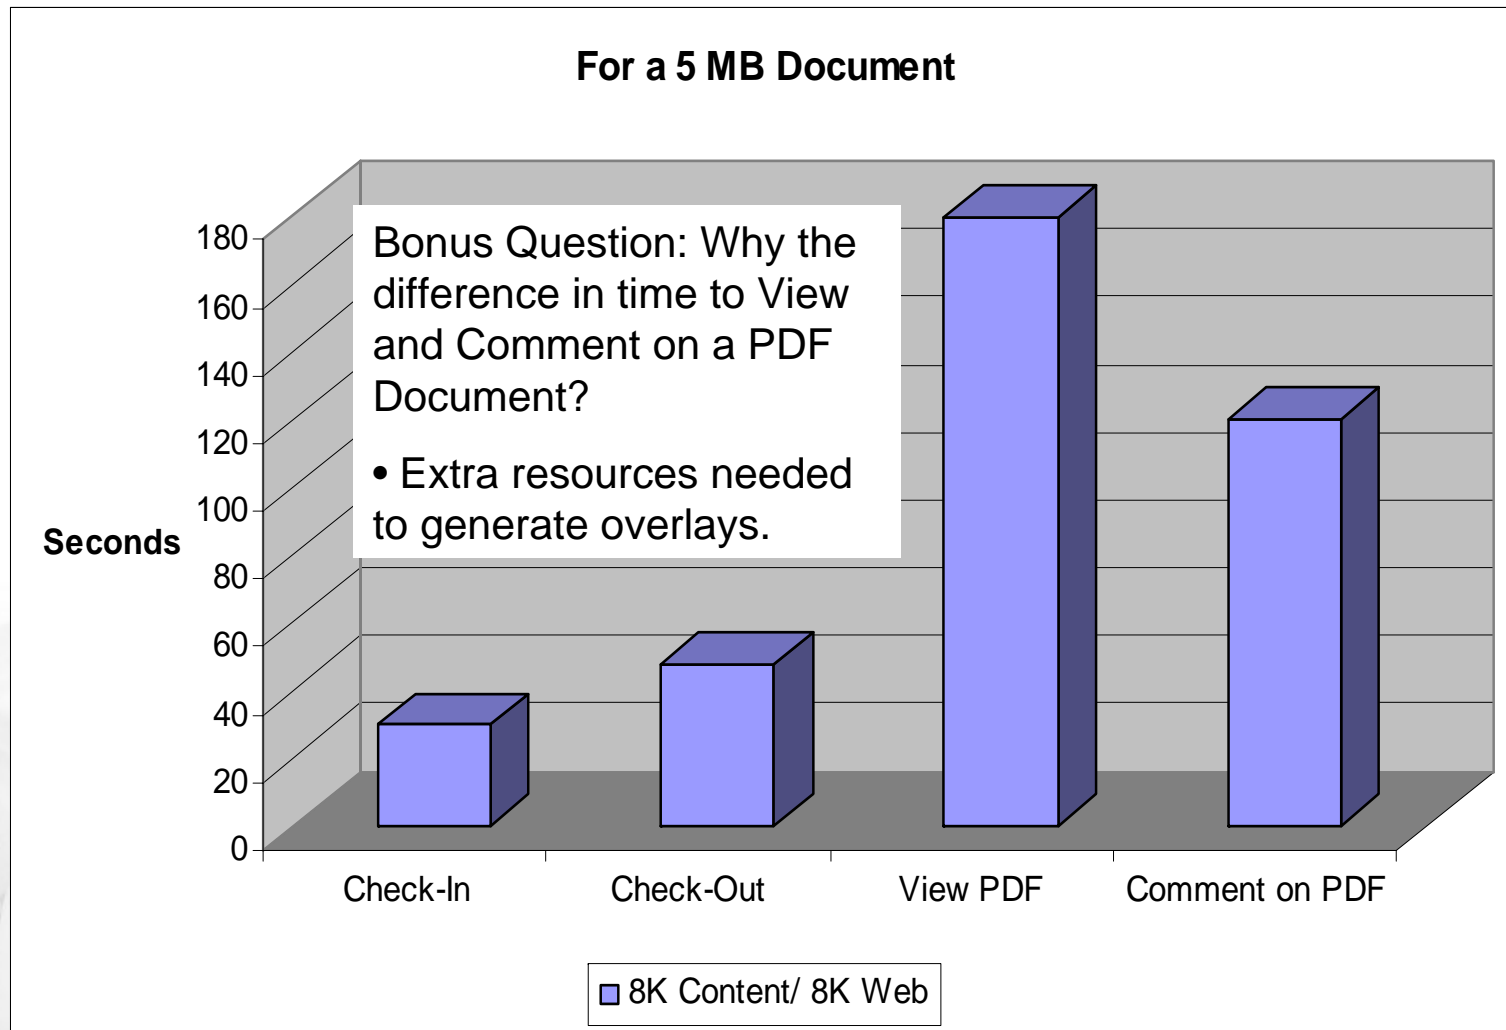


Function vs. Default Window Size





Function vs. Default Window Size





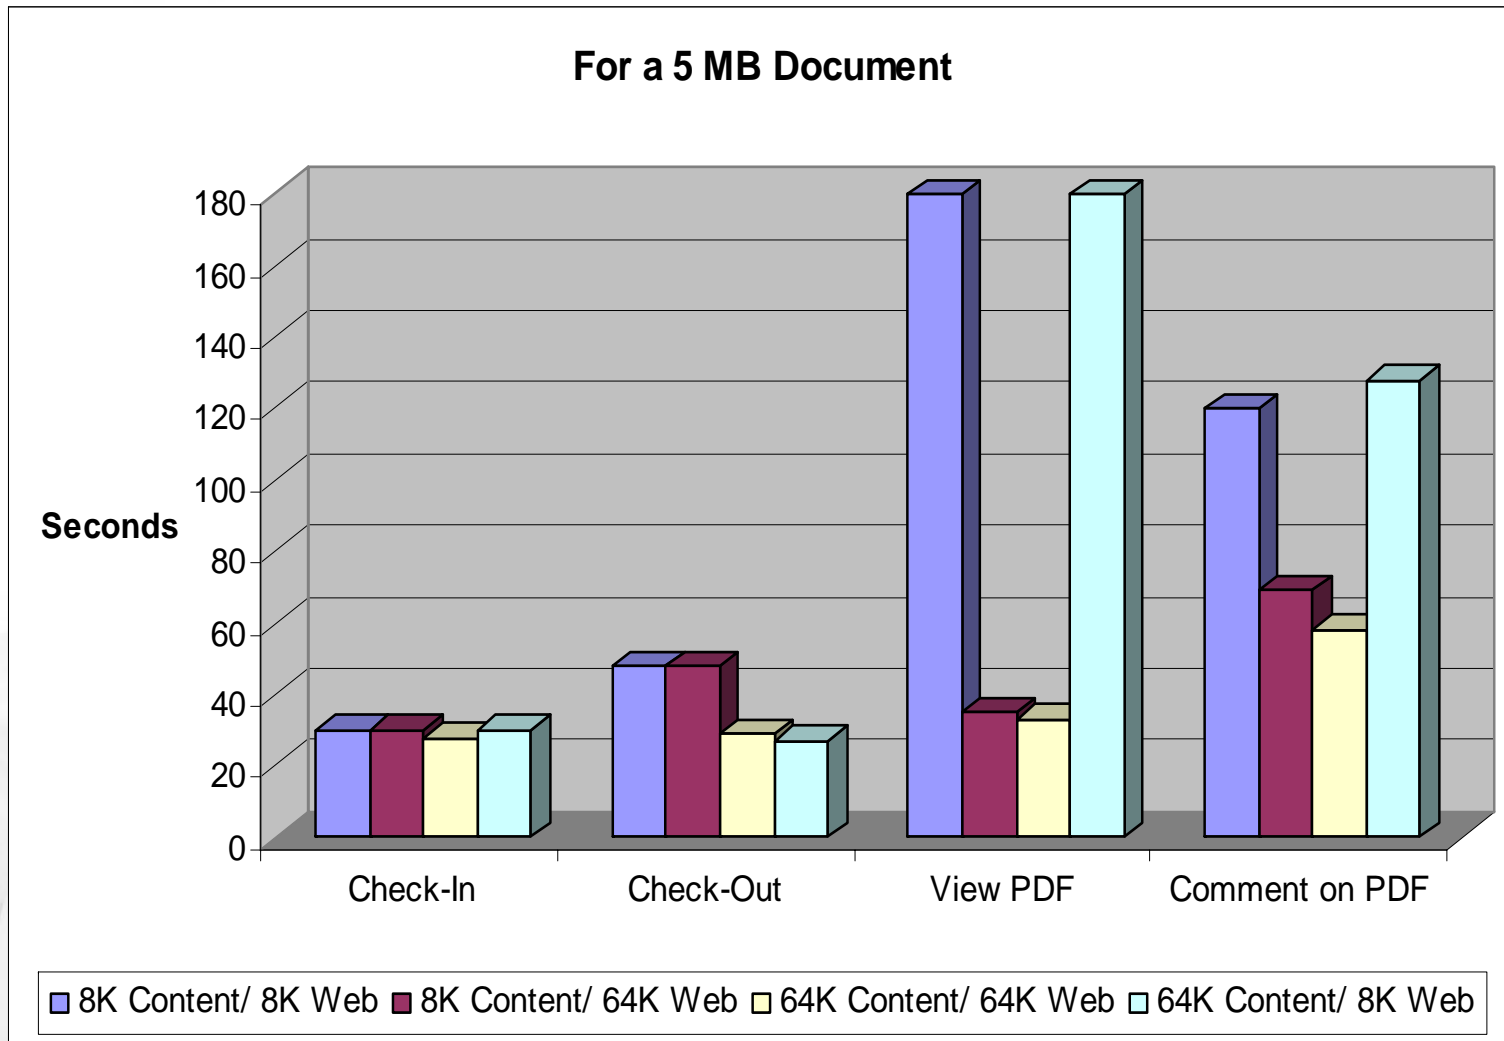
Use The Tools They Give You

- The Lesson 1 (obvious)
 - If the tools are designed to do the work, let them do it
 - We still are not sure why the window size is set in the registry instead of using MSS
- The Lesson 2 (kind of obvious, but took some work to find it)
 - Values in the registry for the TCP/IP Window Size on both the Content and Web server need to be adjusted based on the “typical” size of the documents in the system.
 - The larger the document, the more overhead in creating and assembling the packets.



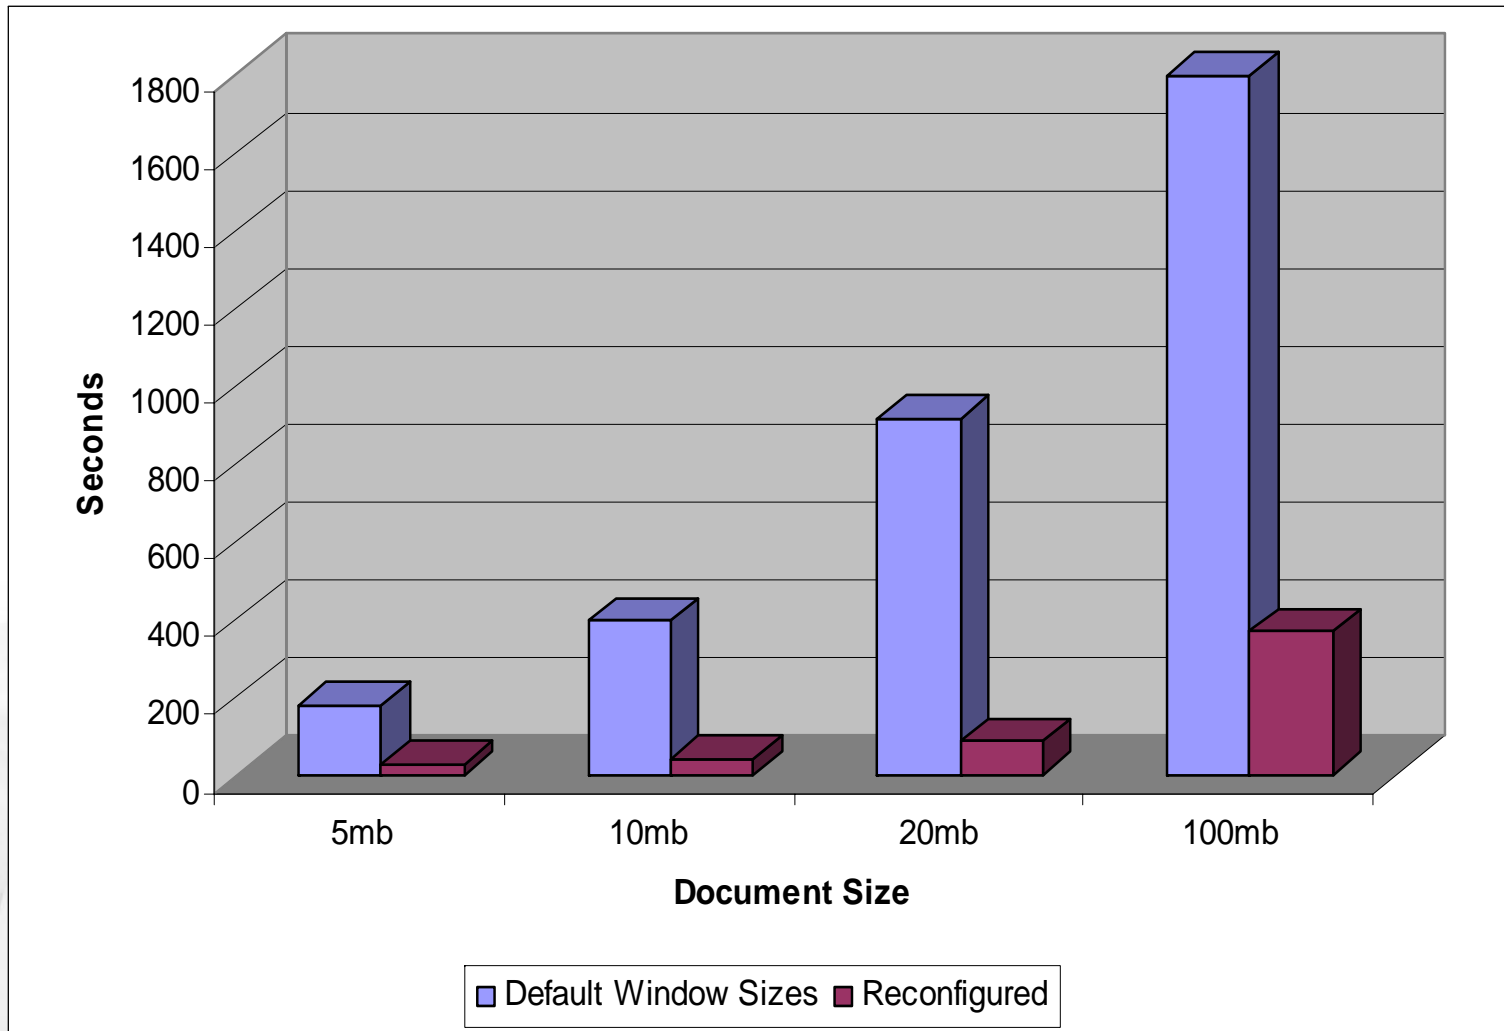


Finding the Best Fit for TCP/IP Window Size





Viewing a PDF Document - Reconfigured





But You Still Have To Turn It On





But You Still Have To Turn It On

- What are the seven jobs every Documentum developer should know and use?
 1. DMClean
 2. DMFilescan
 3. LogPurge
 4. ConsistencyChecker
 5. UpdateStats
 6. QueueMgt
 7. StateOfDocbase





But You Still Have To Turn It On

| | |
|--------------------|---|
| DMClean | Removes deleted and orphaned objects from the Docbase |
| DMFilescan | Removes deleted and orphaned content files from the file system |
| LogPurge | Removes server and session logs from the Docbase and file system |
| ConsistencyChecker | Runs 77 referential integrity checks on the Docbase |
| UpdateStats | Updates database table statistics and repairs fragmented tables |
| QueueMgt | Deletes dequeued Inbox items from the Docbase |
| StateOfDocbase | Produces report of the repository environment and statistics about object types |





But You Still Have To Turn It On

- **Background**
 - After upgrading Eddie, the maintenance jobs were not performed.
 - In approximately 15 months, over 2,000,000 log files accumulated.
 - As part of the weekly reboots, a CHKDSK is performed
 - On one eventful day, CHKDSK took more than 24 hours to complete
 - It was stopped after approx. 36 hours
 - The file system had become corrupted
- Which files do you suspect become corrupted?
 - The log files





But You Still Have To Turn It On

- Lesson 1 (not very obvious)
 - Out of the box, jobs that delete objects or require user customized parameters are not configured to run automatically
- Lesson 2 (should have been very obvious)
 - The need for routine maintenance
 - Lead to the creation of a System Maintenance Manual
 - Lists every maintenance task, when it is to be performed, and by whom





Listen To the Pros





Listen To The Pros

- What is EMC's position with regard to the Content and Index Servers?
 - A. Combining them on one server is acceptable.
 - B. Strong recommendation not to combine them.
 - C. Do not express an position.
- The Answer
 - B. Strong recommendation not to combine them.





Listen To The Pros

- **Background**
 - The most common complaint from users after the Eddie upgrade was “It’s so slow”
 - During the course of the Eddie upgraded an effort was made to contain the budget
 - Not only were the Content and Indexing servers on the same server, the Database server was also on the same server
- **The Lesson (very, very obvious)**
 - CPU utilization was consistently at 100%, and therefore contention for system resources among the various jobs and processes





Listen To The Pros

- What happens when the Index server is disabled (after being combined with the Content server)?
 - A. Nothing
 - B. Impacts _____ functionality.





Listen To The Pros

- **Background**
 - After disabling the Index server (it turns out users were not really searching on the contents of documents), users started reporting that the search results were not returning documents they knew where there.
- **The Lesson (not so obvious)**
 - Turning off the Index server changed the case sensitivity of the Advanced Search.
 - Became case sensitive





Be Careful What You Ask For





Be Careful What You Ask For

- In DCTM 5.3, the STMP calls to send notifications, e.g., promotion failures, are
 - A. Run as a method on the server that can control the number of invocations.
 - B. Made individually through a procedure that executes a .exe





Be Careful What You Ask For

- Background
 - The most common complaint from Eddie users after the upgrade was “It’s so slow”.
 - The Eddie business community also thought it would be worthwhile if everyone in a business are, i.e., Clinical Development, be able to access every document
 - This design also means that when a notification pertaining to a document needs to be sent, then everyone who has access to that document receives the notification
 - The symptom of the problem:
 - Many of the people who should have been receiving the notifications were not receiving them.





Be Careful What You Ask For

- Lesson 1 (very, very obvious)
 - Because system performance was being impacted by multiple servers on the same box, there were not enough system resources to keep up with the demand to send all the notifications.
- Lesson 2 (not obvious at first, but then ...)
 - Be careful what you ask for
 - In trying to please a broad user population, the lack of granularity in the security design lead to:
 - So many notifications, users ignore them all
 - Contention for system resources





Be Careful What You Ask For

- The workaround:
 - Email aliases to groups of individuals on the email system.
- The fix:
 - EMC modified the notification process to queue notifications, rather than trying to send all at once





Don't Take It For Granted...





Don't Take It For Granted...

- True or false, the out-of-box workflow reporting applies the necessary indexes to the tables?
 - A. True
 - B. False





Don't Take It For Granted ...

- **Background**
 - In EDNA, the display of the workflow reporting page was taking about 5 minutes
- **The Fix**
 - Applied indexes to the dmi_queue and dm_audittrail tables increased workflow reporting performance by 15 fold.
 - Page now displays in 20 seconds
- **The Lesson (maybe obvious, maybe not)**
 - Don't take it for granted, that just because you use the out-of-the box configurations that they are optimal
 - Sometimes this may be the first place to look to resolve a problem





It's Never Too Late





It's Never Too Late

- The scenario
 - A document, *imported in PDF format*, is in an approval workflow, and 2 of the 4 approvers have approved the document.
 - The document owner realizes a change is needed to the document, and halts the workflow.
 - The document is revised, a new PDF created and imported, and a new approval workflow is started.
 - All 4 approvers apply their eSignature, and the document is promoted to Approved.
- What happens to the first 2 eSignatures?
 - The original 2 eSignatures are included with the 4 new eSignatures
 - Two signature pages were manifested





It's Never Too Late

- **Background**
 - This was the first instance of such a problem being reported to EMC
 - After significant research, EMC identified this as a “bug”
 - Because of the complexity of the issue, it was determined that a hot fix would not be created
 - A program was developed to find other instances of multiple signature pages in the repository
- **The Lesson (more of a surprise)**
 - Even when getting close to product being “Sunsetted”, new bugs may still be found





The Important Lessons Learned





Important Lessons Learned

- For both Business and Technical personnel
 - Documentum is extremely complex (obvious)
 - Add an application such as DCM, and it gets even more complex (obvious)
 - Add third party products like RenderPerfect or PDF Aqua, and the complexity increases even more (even more obvious)
 - Therefore, finding the cause of and resolving the issues is a complex exercise
 - Extremely obvious
 - But often taken for granted





Important Lessons Learned

- For both Business and Technical personnel
 - It all starts with the design
 - Take your time, prototype where you can
 - Better to spend time getting the design right than addressing issues from a poor design
 - Will also help articulate the requirements
 - Provides insight into what Documentum can and can't do
 - Identifies where changes to business processes need to take place





Important Lessons Learned

- For Business personnel – the users of the system
 - Identifying the source of problems and solving them takes time
 - Be patient
 - Learn about the technical aspects of the system
 - It will help with the communication with the Technical staff
 - It is not always “within our control”
 - Someone knew the version of java being used was “outdated” and installed a newer version on the user’s desktop





Important Lessons Learned

- For Technical personnel
 - Listen to your users
 - Be proactive
 - Act like a user and identify what needs improvement
 - Learn the business processes the system supports
 - This will help understand the issues from the users' perspective
 - Have an inquisitive mind
 - Act like a detective and examine all the clues
 - The obvious solution may not always fix the problem
 - Work with the EMC technical resources
 - Knowledge base
 - Support team





Q & A





Prediction

- Your response to this presentation is:
 - Either
 - Wow, those Eisai guys are smart. They solved some really tough problems.
 - Wow, those Eisai guys really aren't too smart. Why did they do those things to get into those situations?
- But, somewhere out there ...
 - Someone has an even better lesson they have learned.
 - Let's hear it!





Thank You!

