Redbooks
ibm.com/redbooks

IBM Software for SAP Solutions

Yaro Dunchych
Peter Bahrs
Khirallah Birkler
Bernd Eberhardt
Navneet Goyal
James Hunter
Derek Jennings
Joe Kaczmarek
Michel Laaroussi
Michael Love
Stefan Momma
Nick Norris
Martin Oberhofer
Manfred Oevers
Paul Pacholski
Andrew Stalnecker
Jörg Stolzenberg
Pierre Valiquette

ECM

IBM

Redbooks

International Technical Support Organization

IBM Software for SAP Solutions

September 2015

SG24-8230-01

Note: Before using this information and the product it supports, read the information in “Notices” on page ix.

Second Edition (September 2015)

This edition applies to Version 2, Release 0, Modification 0 of the IBM Reference Architecture for SAP.

© Copyright International Business Machines Corporation 2014, 2015. All rights reserved.
Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Notices ix
Trademarks x

IBM Redbooks promotions xi

Preface xiii
Authors xiv
Now you can become a published author, too! xix
Comments welcome xx
Stay connected to IBM Redbooks xx

Summary of changes xxi
September 2015, Second Edition xxi

Chapter 1. Why IBM software matters in SAP solutions 1
1.1 Overview 2
1.2 Critical success factors for an SAP-centric transformation 2
1.2.1 Deploying a system of engagement for SAP 4
1.2.2 Balancing SAP with an application-independent, industry-leading integration platform solution 4
1.2.3 Establishing governance for architectural decisions 5
1.2.4 Avoiding custom coding 5
1.3 Combined value of IBM and SAP software 6
1.3.1 Reduced business and IT risk 6
1.3.2 Accelerated SAP integration into a heterogeneous enterprise 7
1.3.3 Business agility 7
1.3.4 Cost reduction 8

Chapter 2. IBM Reference Architecture for SAP 9
2.1 Overview 10
2.2 Architecture goals 10
2.2.1 Use standard, non-customized SAP applications 10
2.2.2 Reuse pre-built SAP integration 11
2.2.3 Use best-in-class technologies when extending beyond the SAP domain 11
2.2.4 Use open, well-established standards 11
2.2.5 Use pre-built software capabilities provided by IBM 11
2.3 IBM Reference Architecture for SAP overview 12
2.3.1 Systems of engagement, record, and interaction 12
2.3.2 Services view 14
2.3.3 Application integration: Inner ring and outer ring architecture 16
2.3.4 Enterprise integration services 19
2.3.5 Process optimization 21
2.3.6 User interface 24
2.3.7 Master data management (MDM) 29
2.3.8 Enterprise content management 31
2.3.9 Business analytics 32
2.3.10 DevOps for SAP 34

Chapter 3. Enterprise integration services for SAP 39

© Copyright IBM Corp. 2014, 2015. All rights reserved.

3.1 Introduction to enterprise integration services for SAP applications 40
3.2 Architecture goals 41
3.2.1 Align enterprise integration services with SAP implementation methodology. . . . 41
3.2.2 Use best-in-class technologies for custom integration development 41
3.2.3 Minimize costs of integration for non-strategic systems 42
3.2.4 Loosely coupled applications 42
3.2.5 Use open, well-established standards 42
3.3 Scenarios and patterns for ongoing integration with SAP 42
3.3.1 Identifying integration scenarios 43
3.3.2 Common integration patterns 48
3.4 Architecture overview of ongoing integration with SAP 51
3.5 Architecture components of ongoing integration with SAP 52
3.5.1 Enterprise Service Bus 52
3.5.2 Extract, transform, and load (ETL) 55
3.5.3 Service governance 57
3.5.4 Reliable File Transfer 60
3.5.5 Process services 63
3.5.6 Logging and error handling 64
3.5.7 Integration workload placement guidelines: ESB versus ETL. 64
3.6 Initial data load 66
3.7 References 73

Chapter 4. Process optimization for SAP 75
4.1 SAP solutions as a system of engagement. 76
4.2 Architecture overview 77
4.3 SAP active business performance optimization architecture 78
4.4 IBM Smarter Process for SAP capabilities 81
4.4.1 SAP Solution Manager integration 83
4.4.2 SAP Guided Workflow 86
4.4.3 Process orchestration, integration, and event management. 91
4.4.4 Process discovery and monitoring 95
4.4.5 Iterative business blueprinting 100
4.4.6 Decision automation 102
4.4.7 Process automation 103
4.5 IBM Smarter Process for SAP products and solutions 104
4.6 How IBM Smarter Process for SAP creates sustained business value. 105
4.7 IBM Smarter Process for SAP usage scenarios 106
4.7.1 IBM Smarter Process for SAP in the phases of an SAP project 107
4.7.2 Post-implementation value augmentation 112
4.8 Conclusion 114
4.8.1 Capability and value summary 114
4.8.2 IBM Smarter Process for SAP Affinity Analysis and Business Value Assessment Workshop 114
4.9 Other IBM Software Group publications, assets, and tools. 116
4.10 IBM Global Business Services SAP assets and tools 116
4.11 References 116

Chapter 5. Mobile access for SAP 117
5.1 IBM MobileFirst overview 118
5.2 Spectrum of mobile app development approaches 119
5.3 IBM MobileFirst for SAP architectures 121
5.3.1 Architecture goals for SAP mobile enablement in a heterogeneous enterprise . 121
5.3.2 IBM MobileFirst for SAP architecture overview 124

5.3.3 Fast-track SAP mobile enablement with IBM Worklight and SAP NetWeaver Gateway . . . 125
5.3.4 IBM MobileFirst integration with SAP with no moving parts . . . 129
5.3.5 Accelerated mobile integration with SAP using IBM WebSphere Cast Iron . . . 129
5.3.6 Full featured mobile integration with SAP using IBM Integration Bus . . . 132
5.3.7 Access to existing SAP Fiori Apps using IBM MaaS360 . . . 134
5.4 Optional components driving enhanced features in mobile architectures . . . 136
5.4.1 Enhancing mobile architectures by adding IBM API Management capabilities . 136
5.4.2 Enhancing mobile architectures by adding IBM mobile analytics and quality assurance capabilities . . . 138
5.4.3 Enhancing mobile architectures by adding secure offline capabilities . . . 139
5.5 Lessons learned from actual projects . . . 140
5.5.1 Direct connectivity from mobile applications to SAP is rarely used . . . 140
5.5.2 Late decision on native versus hybrid apps . . . 141
5.5.3 Adding mobile business analytics features dynamically . . . 141
5.5.4 Separation of security domains . . . 141
5.6 References . . . 142

Chapter 6. Portal integration with SAP . . . 143
6.1 Overview of integrating IBM WebSphere Portal with SAP applications . . . 144
6.2 Architecture overview . . . 144
6.3 Types of use cases . . . 146
6.3.1 Casual use cases . . . 146
6.3.2 Detailed use cases . . . 146
6.4 WebSphere Portal integration with SAP app use cases . . . 147
6.4.1 Federated portal . . . 147
6.4.2 Integrating with the web application bridge feature . . . 147
6.4.3 Integrating with IBM WebSphere Portal Integrator for SAP . . . 150
6.4.4 Integrating with Web Services for Remote Portlets (WSRP) . . . 152
6.5 Service-level integration . . . 153
6.5.1 Direct integration with SAP applications using SAP Java connector . . . 154
6.5.2 Integrating with an enterprise service bus to connect to SAP applications . . . 155
6.5.3 Integrating with SAP NetWeaver Gateway . . . 155
6.6 Architecture guidelines . . . 156
6.7 Summary . . . 157
6.8 References . . . 157

Chapter 7. Master data management for SAP . . . 159
7.1 Master data management introduction . . . 160
7.2 Why master data management is important for SAP applications . . . 161
7.3 Overview of IBM Master Data Management capabilities . . . 163
7.4 Architecture goals . . . 166
7.5 Architecture overview . . . 166
7.6 IBM InfoSphere MDM for SAP applications . . . 168
7.7 Architecture patterns . . . 170
7.7.1 Master Data Integration . . . 172
7.7.2 Master data distribution . . . 175
7.7.3 MDM hub patterns and MDM implementation styles . . . 179
7.7.4 Selecting MDM hub and MDM implementation styles for environments with SAP applications . . . 185
7.8 References . . . 187

Chapter 8. Enterprise Content Management for SAP . . . 189
8.1 Enterprise content management business goals . . . 190

8.1.1 Information lifecycle management: More than just archiving 190
8.1.2 Information lifecycle governance applied to SAP systems 194
8.1.3 IBM proposes a base structure of an integrated ECM solution. 195
8.2 ECM for SAP use cases and solution architecture 196
8.2.1 SAP archiving standards. 196
8.2.2 SAP archiving use cases 197
8.2.3 Architecture of IBM Content Collector for SAP Applications. 200
8.2.4 IBM Datacap 206
8.3 Business process enhancements through ECM for SAP solutions. 207
8.3.1 Objectives of a document-oriented workflow management 207
8.3.2 SAP-centric versus ECM-centric process management 208
8.3.3 Components of an ECM for SAP Solution 209
8.3.4 Capturing solution components 210
8.3.5 ECM SAP solution architecture. 220
8.4 Data governance: Managing growth and compliance 221
8.4.1 Business drivers 221
8.4.2 SAP infrastructure for data archiving 224
8.4.3 Data archiving and the choice of IBM ECM content repositories 225
8.4.4 SAP ArchiveLink-based data archiving. 226
8.4.5 Data archiving using SAP ILM 228
8.4.6 Comparison of SAP ArchiveLink and ILM-based data archiving. 229
8.4.7 Adding the value of IBM middleware and storage solutions for SAP data archiving purposes 229
8.5 References 230

Chapter 9. IBM Business Analytics infrastructure for SAP. 231
9.1 IBM Business Analytics infrastructure for SAP value proposition 232
9.1.1 Architecture overview 235
9.2 IBM Business Analytics integration architectures 237
9.2.1 IBM Enterprise Data Warehouse products 238
9.2.2 IBM InfoSphere DataStage 239
9.2.3 IBM InfoSphere Data Replication 239
9.3 Detailed review of IBM Business Analytics integration architectures for SAP. 239
9.3.1 Data export from SAP Business Suite into an IBM enterprise data warehouse . 239
9.3.2 Data export from SAP BW into an IBM EDW 240
9.3.3 Operational analytics with Cognos Business Intelligence directly accessing SAP solutions 242
9.3.4 Managing business performance with SAP and IBM Cognos TM1 244
9.3.5 Predictive analytics with SAP 245
9.4 Conclusion 247
9.5 References 248

Chapter 10. DevOps for SAP 249
10.1 IBM DevOps for SAP overview 250
10.1.1 IBM DevOps for SAP key capabilities. 251
10.2 Application lifecycle management for SAP 253
10.2.1 IBM Rational Connector for SAP Solution Manager 254
10.2.2 Requirements management 254
10.2.3 Blueprint push from SAP Solution Manager 256
10.2.4 Project planning and execution. 257
10.2.5 Change, defect, and incident management 259
10.2.6 Quality management. 260
10.2.7 Impact analysis 261

10.3 Collaborative development for SAP 262
10.3.1 Improve SAP developer and team productivity 264
10.3.2 Real-time visibility into SAP Agile delivery and maintenance projects 265
10.3.3 Accelerate agile development adoption and results 266
10.4 Continuous testing for SAP 266
10.4.1 Functional testing 266
10.4.2 Integration testing and service virtualization 267
10.4.3 Performance testing 268
10.5 Continuous release and deployment for SAP 268
10.6 Continuous business planning for SAP 269
10.6.1 Enterprise architecture 270
10.7 Summary 272
10.8 References 273

Chapter 11. Systems security for SAP 275
11.1 SAP systems and IBM security management integration overview 276
11.1.1 Key capabilities: Logical components of security reference architecture 277
11.1.2 Mapping logical security architecture components to IBM and SAP software . 278
11.2 SAP systems security and IBM security management integration scenarios 279
11.2.1 Key solution components 279
11.2.2 Generic components and concepts of a security architecture 280
11.3 Identity system scenarios 281
11.3.1 Identity management scenario 281
11.3.2 Identity feed scenario 283
11.3.3 User provisioning scenario 284
11.4 Authentication system scenarios 285
11.4.1 Access management scenario 285
11.4.2 Single sign-on (SSO) scenario 287
11.4.3 Identity propagation scenario 289
11.5 Authorization system scenario 291
11.6 Audit system scenarios 293
11.6.1 Security monitoring and analytics scenario 294
11.6.2 Source code analysis scenario 296
11.7 Identity management products and solutions 297
11.7.1 IBM Security Identity Manager 297
11.7.2 IBM Security Directory Integrator 297
11.7.3 IBM Security Directory Server 298
11.8 Access management products and solutions 299
11.8.1 IBM Security Access Manager for Enterprise Single Sign-On 299
11.8.2 IBM Security Access Manager for Web 299
11.8.3 IBM Tivoli Federated Identity Manager 300
11.9 Audit products and solutions 302
11.9.1 IBM InfoSphere Guardium 302
11.9.2 IBM Security QRadar Log Manager 302
11.9.3 IBM Security QRadar Risk Manager 303
11.9.4 IBM Security QRadar SIEM 303
11.9.5 IBM Security AppScan 304
11.10 References 305

Chapter 12. Systems management for SAP 307
12.1 Architectural goals 308
12.1.1 Enable optimal availability and usability of complex business systems 308
12.1.2 Provide visibility to unplanned business process outages 308

12.1.3 Enable historical view of business process availability 309
12.2 Business process availability management overview 309
12.2.1 Complex IT solutions require multiple levels of systems management. 309
12.2.2 Multiple systems management tools exist for each layer of solution 310
12.2.3 Systems management considerations 311
12.3 Systems management reference architecture for SAP-centric solutions 312
12.3.1 Application architecture. 312
12.3.2 Infrastructure architecture 313
12.3.3 Systems management architecture. 313
12.4 Business process availability management for SAP-centric solutions 315
12.4.1 Business process availability management architecture. 316
12.4.2 Business Process DLA overview 318
12.5 Summary 320
12.6 References 321

Related publications 323
IBM Redbooks 323
Other publications 323
Online resources 324
Help from IBM 324

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features described in this document in other countries. Consult your local IBM representative for information about the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:
IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not enable disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

© Copyright IBM Corp. 2014, 2015. All rights reserved.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at http://www.ibm.com/legal/copytrade.shtml

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

AppScan®
Cast Iron®
CICS®
Cognos®
Concert™
Daeja™
DataPower®
DataStage®
DB2®
developerWorks®
DOORS®
FileNet®
Global Business Services®
Guardium®
IBM®
IBM MobileFirst™
IBM SmartCloud®
IBM UrbanCode™
IBM Watson™
IMS™
InfoSphere®
NetView®
OMEGAMON®
Optim™
POWER7®
PureApplication®
PureData®
QRadar®
QualityStage®
Rational®
Rational Team Concert™
Redbooks®
Redpapers™
Redbooks (logo) ®
SPSS®
System z®
System z10®
Tealeaf®
Tivoli®
TM1®
WebSphere®
Worklight®
z/OS®
z10™
zEnterprise®

The following terms are trademarks of other companies:

Adobe, the Adobe logo, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

MaaS360, Secure Productivity Suite, and We do IT in the Cloud. device are trademarks or registered trademarks of Fiberlink Communications Corporation, an IBM Company.

Connect:Direct, Netezza, and N logo are trademarks or registered trademarks of IBM International Group B.V., an IBM Company.

Microsoft, Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Java, and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Other company, product, or service names may be trademarks or service marks of others.

Find and read thousands of IBM Redbooks publications

- Search, bookmark, save and organize favorites
- Get up-to-the-minute Redbooks news and announcements
- Link to the latest Redbooks blogs and videos

Get the latest version of the Redbooks Mobile App

iOS **Download Now** Android

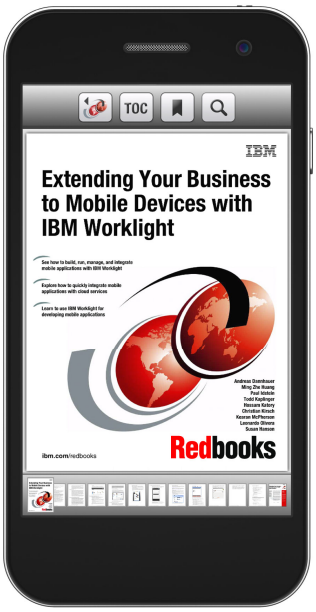

Promote your business in an IBM Redbooks publication

Place a Sponsorship Promotion in an IBM® Redbooks® publication, featuring your business or solution with a link to your web site.

Qualified IBM Business Partners may place a full page promotion in the most popular Redbooks publications. Imagine the power of being seen by users who download millions of Redbooks publications each year!

ibm.com/Redbooks
About Redbooks → Business Partner Programs

THIS PAGE INTENTIONALLY LEFT BLANK

Preface

SAP is a market leader in enterprise business application software. SAP solutions provide a rich set of composable application modules, and configurable functional capabilities that are expected from a comprehensive enterprise business application software suite.

In most cases, companies that adopt SAP software remain heterogeneous enterprises running both SAP and non-SAP systems to support their business processes. Regardless of the specific scenario, in heterogeneous enterprises most SAP implementations must be integrated with a variety of non-SAP enterprise systems:

- Portals
- Messaging infrastructure
- Business process management (BPM) tools
- Enterprise content management (ECM) methods and tools
- Business analytics (BA) and business intelligence (BI) technologies
- Security
- Systems of record
- Systems of engagement

The tooling included with SAP software addresses many needs for creating SAP-centric environments. However, the classic approach to implementing SAP functionality generally leaves the business with a rigid solution that is difficult and expensive to change and enhance.

When SAP software is used in a large, heterogeneous enterprise environment, SAP clients face the dilemma of selecting the correct set of tools and platforms to implement SAP functionality, and to integrate the SAP solutions with non-SAP systems.

This IBM® Redbooks® publication explains the value of integrating IBM software with SAP solutions. It describes how to enhance and extend pre-built capabilities in SAP software with best-in-class IBM enterprise software, enabling clients to maximize return on investment (ROI) in their SAP investment and achieve a balanced enterprise architecture approach. This book describes IBM Reference Architecture for SAP, a prescriptive blueprint for using IBM software in SAP solutions. The *reference architecture* is focused on defining the use of IBM software with SAP, and is not intended to address the internal aspects of SAP components.

The chapters of this book provide a specific reference architecture for many of the architectural domains that are each important for a large enterprise to establish common strategy, efficiency, and balance. The majority of the most important architectural domain topics, such as integration, process optimization, master data management, mobile access, enterprise content management, business intelligence, DevOps, security, systems monitoring, and so on, are covered in the book.

However, several other architectural domains exist that are not described in the book. This is not to imply that these other architectural domains are not important or are less important, or that IBM does not offer a solution to address them. It is reflective only of time constraints, available resources, and the complexity of assembling a book on an extremely broad topic.

Although more content could have been added, the authors are confident that the scope of architectural material that has been included should provide organizations with a strong head start in defining their own enterprise reference architecture for many of the important architectural domains, and it is hoped that this book provides great value to those reading it.

© Copyright IBM Corp. 2014, 2015. All rights reserved.

This book is targeted to the following audiences:

- Client decision makers and solution architects leading enterprise transformation projects and wanting to gain further insight so that they can benefit from the integration of IBM software in large-scale SAP projects.
- IT architects and consultants integrating IBM technology with SAP solutions.

Authors

This book was produced by a team of specialists from around the world working with the IBM International Technical Support Organization (ITSO).

Yaro Dunchych is an Executive Architect in IBM Software Services. Yaro leads IBM software deployments at IBM clients, including internal adoption of IBM software in the IBM company in such areas as smarter process, mobile, integration, and cloud technologies. Yaro leads delivery of IBM software solutions for SAP based on the IBM Reference Architecture for SAP, a large-scale asset that provides a comprehensive blueprint of IBM software capabilities for SAP. During his 15 years with IBM, Yaro created multiple cross-brand solutions from services engagements that led to several innovations. Yaro holds a Ph.D. degree in Electronic Engineering from Air Force Engineering Academy in Moscow, Russia.

Peter Bahrs is a Distinguished Engineer and Chief Technology Officer (CTO) for IBM WebSphere® and Software Lab Services. Peter is focused on mobile, connectivity and integration, and Internet of Things (IoT) technologies. Peter leads projects in banking, retail, defense, smarter transportation, and insurance. Peter has been the technical executive leader in IBM Software Group (SWG) focusing on integrating SAP solutions with IBM software. Peter has B.S., M.S., and Ph.D. degrees in Computer Science from the Center for Advanced Computer Studies at the University of Louisiana.

Khirallah Birkler is a Certified IT Specialist and IBM/SAP Mobile Solution Architect in the IBM SWG division, IBM Germany. Khirallah helps clients to identify requirements, design, and architect mobile solutions based on the IBM MobileFirst™ portfolio. A special focus of his work is the integration of enterprise data into mobile scenarios, which is always a key requirement for enterprise customers. Khirallah holds a Bachelor of Science degree in Information Technology Management from the University of Cooperative Education, Stuttgart.

Bernd Eberhardt is a Product Manager at the IBM SAP International Competence Center (ISICC) in Walldorf, Germany. Bernd is the product manager for the IBM Rational® SAP Alliance. During his 14-year tenure with IBM, Bernd had various assignments in sales and consulting services, with a strong focus on IBM Rational quality management solutions.

Navneet Goyal is a Senior Developer in the IBM Software Lab organization in IBM India. Navneet designs, develops, and supports the IBM web application bridge feature for the IBM WebSphere Portal and business portlets of the IBM WebSphere Portal product. He holds a Bachelor of Engineering degree in Computer Engineering from the Delhi College of Engineering, Delhi, India.

James Hunter is a Consulting, System Integration, and SAP Solution Offering Leader in IBM Rational organization, IBM SWG in the UK. He spent a large part of his early career architecting and implementing critical control, large-scale, distributed systems in the finance, public, and defense sectors. His university studies were on politics and economics with computer science, and he is a Chartered Fellow of the British Computer Society.

Derek Jennings is a Senior Certified Consulting IT Specialist with the IBM Global Business Services® division in the US. Derek is currently an offerings and solutions architect for IBM Dev/Test Cloud Services. Derek has over twenty years of experience in full lifecycle performance engineering for SAP and large, complex enterprise systems. He has also designed the monitoring and management strategy for many of the largest and most mission-critical business systems in IBM.

Joe Kaczmarek is a Global Executive for IBM Smarter Process for SAP. Joe leads IBM Smarter Process for SAP business for IBM worldwide. Over the past nine years, he has held numerous global sales and business unit executive roles within IBM WebSphere Application Integration and Middleware brand in IBM SWG. Before joining IBM SWG, Joe was a consulting practice leader in IBM Global Business Services and PricewaterhouseCoopers. Joe also spent 15 years as an entrepreneur founding, leading, and harvesting IT and marketing services companies.

Michel Laaroussi is an SAP BA consultant with IBM Canada Global Business Services. He helps clients by providing them with agile SAP BA strategy roadmaps and implementing SAP BA lifecycle projects based on the SAP Business Analytics portfolio (SAP Business Warehouse, SAP High-Performance Analytic Appliance (HANA), and SAP Business Objects). Michel advises clients on enterprise architecture guidelines aligned with their corporate and business strategies, and with SAP's product strategies. Michel holds a Bachelor of Science degree from Université Paris Diderot, Paris 7 and Université Pierre et Marie Curie, Paris 6, and a Master's degree in Competitive Intelligence from Economic Warfare School of Paris (EGE ESLSCA).

Michael Love is an industry leader with a mission to help organizations worldwide to adopt a superior methodology of integrating packaged applications across the enterprise, therefore significantly reducing costs and enabling agility for the business. In his current Executive Architect role, Michael is responsible for working with IBM customers who also have a significant investment in SAP applications. Michael's focus is in practical solutions to help customers lower existing integration costs by incorporating an enterprise integration strategy using SOA methodologies. Michael has the experience of directly engaging with over 300 of the largest SAP customers worldwide to guide them through integration, business process management, and enterprise architecture leading practices relative to balancing their SAP environment.

Stefan Momma is a Software Architect in the Enterprise Content Management division of IBMSWG at the IBM Research and Development Lab in Boeblingen, Germany. He is responsible for the product architecture of IBM ECM for SAP solutions. Before his current assignment, Stefan focused on full-text search in IBM DB2® and IBM Enterprise Content Management portfolio products. He holds a Diploma degree in Computer Science and Computational Linguistics from the University of Stuttgart.

Nick Norris is an IBM certified Executive IT Specialist currently working as IBM Rational Worldwide Lead Architect for the Financial Services Sector and SAP solutions. In this role, Nick works with many organizations around the world to understand their needs, determine which combination of IBM Rational solutions are the best fit for them, and how they might be best used to deliver both near-term and long-term value. Nick is a frequent speaker at technology conferences. He authored a series of articles on *Governing and managing enterprise models* in IBM developerWorks®, and the article series *Software Development Compliance - Work Authorization and Requirements Integrity* in jazz.net. He is the co-author of the IBM Redpapers™ publication *Building Service-Oriented Banking Solutions with IBM Banking Industry Models and Rational SDP*, REDP-4232.

Martin Oberhofer is an Executive Architect in Information Management development, IBM SWG. Martin works on Enterprise Information Architecture with large clients worldwide. He helps customers define their enterprise information strategy and architecture, solving information-intense business problems. Martin co-authored the books *Enterprise Master Data Management: An SOA Approach to Managing Core Information* and *The Art of Enterprise Information Architecture: A Systems-Based Approach for Unlocking Business Insight*, in addition to numerous IBM Research articles and IBM developerWorks articles. As an inventor, he contributed to over 70 patent applications for IBM, achieving the IBM Master Inventor title. Martin is also a The Open Group Master Distinguished Architect, and holds a Master's degree in Mathematics from the University of Constance, Germany.

Manfred Oevers is a Development Manager with the IBM ECM division of IBM Software Group, in IBM Germany. Manfred has 14 years of experience in the IT industry. He has worked as a high-performance computing specialist, emerging technology services, life science development, WebSphere software development, ECM Lab services, and release manager. He holds a degree in Physics from the University of Bielefeld, and a Ph.D. in Theoretical Particle Physics from the same university.

Paul Pacholski After an 18-year career as a Senior Developer in the IBM Development Lab in Toronto, Paul joined the Worldwide WebSphere Technical Sales Team in 1999. As a WebSphere Smarter Process Technical Sales Lead, Paul is responsible for the technical enablement of the IBM Smarter Process Technical Sales organization worldwide. Paul's responsibilities also include customer engagements, consulting, presenting at technical conferences, and publishing technical papers. In his most recent role as a Lead Architect, Paul is leading a team that develops SAP capabilities in IBM Smarter Process.

Andrew Stalnecker is a Technical Sales Consultant with the IBM Sales & Distribution division in North America. For the past 14 years, he has helped clients identify requirements, and design and architect solutions based on the IBM ECM portfolio. He is known as the “go-to” person for SAP ECM integration.

Jörg Stolzenberg is the Product Manager for IBM Content Collector and for ECM SAP integrations. After spending 12 years in the ECM industry, mostly in technical positions, Jörg joined IBM FileNet® in 2003 as a Technical Alliance Manager for SAP, based in Walldorf, Germany. From 2005 on, he was also in charge of the Product Management of FileNet’s SAP connectors. Since the acquisition of FileNet by IBM, Jörg is the Product Manager for IBM ECM for SAP product portfolio. From 2009 to 2012 he was also in charge of IBM Content Classification. Since 2012, Jörg’s responsibilities also include IBM Content Collector. Jörg holds a Doctoral degree in Mechanical Engineering.

Pierre Valiquette is a Product Manager for IBM Business Analytics, specializing in SAP Business Warehouse (BW) and SAP HANA connectivity. He has held customer-facing roles throughout his 17-year tenure at IBM, including customer support, software engineer, and third-level support manager. For the past seven years, Pierre has been leading the SAP BW and SAP HANA investments from an IBM development and product management perspective. Pierre continues to work directly with customers, partnering with them to ensure that IBM delivers the wanted functionality moving forward.

The project that produced this publication was managed by **Marcela Adan**, IBM Redbooks Project Leader with ITSO, Global Content Services.

Thanks to the following people for their contributions to this project:

Stefan Liesche, Sascha Schefenacker
IBM Collaboration Solutions Development, IBM Software Group

Jon Deng, Albert Maier, Dylan Murphy, Dan Wolfson
Information Management, IBM Software Group

Scott Schumacher
IBM InfoSphere®, IBM Software Group

Martha Miller
Business Analytics, IBM Software Group

Dave Tropeano
IBM Rational

Jane Hendricks
IBM Business Analytics and Predictive Analytics, IBM Software Group

Aleksandr Nartovich
Worldwide sales, IBM Software Group

Holger Martin, Dorothee Stork
IBM Enterprise Content Manager Technical Sales, IBM Germany

Carsten Steck
SAP Archiving Solutions, IBM Global Business Services

Gang Chen, Sean Sundberg
IBM Software Services for WebSphere

Ingo Dressler
IBM Systems & Technology Group

David Moore
IBM Security Systems, IBM Software Group

Michael Campbell, Robert Kennedy
IBM Security Sales Enablement, IBM Software Group

Greg Truty
IBM MobileFirst Platform, IBM Software Group

Julianne Bielski, Mathew Davis
Cloud & Smarter Infrastructure, IBM US

Volker Kohlstetter, Uta Beyer
xft GmbH, Walldorf, Germany

Now you can become a published author, too!

Here's an opportunity to spotlight your skills, grow your career, and become a published author—all at the same time! Join an ITSO residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run two - six weeks in length, and you can participate either in person or as a remote resident working from your home base.

Learn more about the residency program, browse the residency index, and apply online:

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

We want our books to be as helpful as possible. Send us your comments about this book or other IBM Redbooks publications in one of the following ways:

- Use the online **Contact us** review Redbooks form:

  ibm.com/redbooks

- Send your comments in an email:

  redbooks@us.ibm.com

- Mail your comments:

  IBM Corporation, International Technical Support Organization
  Dept. HYTD Mail Station P099
  2455 South Road
  Poughkeepsie, NY 12601-5400

Stay connected to IBM Redbooks

- Find us on Facebook:

  http://www.facebook.com/IBMRedbooks

- Follow us on Twitter:

  http://twitter.com/ibmredbooks

- Look for us on LinkedIn:

  http://www.linkedin.com/groups?home=&gid=2130806

- Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:

  https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm

- Stay current on recent Redbooks publications with RSS Feeds:

  http://www.redbooks.ibm.com/rss.html

Summary of changes

This section describes the technical changes made in this edition of the book. This edition might also include minor corrections and editorial changes that are not identified.

Summary of Changes
for SG24-8230-01
for IBM Software for SAP Solutions
as created or updated on September 29, 2015.

September 2015, Second Edition

This revision includes the following changed information.

- Section 8.2, “ECM for SAP use cases and solution architecture” on page 196 is rewritten to reflect the changes in the client infrastructure of IBM Content Collector for SAP Applications introduced with version 4.0, which is now based on IBM Content Navigator.

Note: The remaining content of this book has not been updated since the original version published in November 2014. The unchanged information reflects the product names and references that were current in November 2014.

© Copyright IBM Corp. 2014, 2015. All rights reserved.

1

Why IBM software matters in SAP solutions

SAP is a well-established leader in the packaged applications market. The term *packaged application* typically refers to upscale enterprise software suites, such as enterprise resource planning (ERP) or customer relationship management (CRM).

SAP software provides a rich set of composable application modules and configurable functional capabilities that are expected from a comprehensive enterprise application software suite. However, in most cases, organizations that adopt SAP software still remain heterogeneous enterprises. In heterogeneous enterprises, most SAP implementations must be integrated with a variety of non-SAP enterprise systems, portals, messaging infrastructure, security, systems of record, systems of engagement, and more.

The purpose of this chapter is to explain the value of using IBM software in SAP solutions. It reviews a set of critical success factors in large-scale SAP-centric enterprise transformation projects, and explains how IBM software increases the value of SAP-centric business transformation programs.

This chapter includes the following topics:

- 1.1, “Overview” on page 2
- 1.2, “Critical success factors for an SAP-centric transformation” on page 2
- 1.3, “Combined value of IBM and SAP software” on page 6

© Copyright IBM Corp. 2014, 2015. All rights reserved.

1.1 Overview

SAP is a market leader in enterprise business application software. SAP software, although ready-made, rarely comes ready to run. SAP adoption typically becomes an enterprise business transformation program. The SAP adoption typically takes six months to 10 years and the average life span is 5 - 15 years.

Organizations that adopt SAP software keep using both SAP and non-SAP systems to support their business processes. Even in a large-scale SAP adoption, a significant part of the business might continue to use non-SAP enterprise applications for various reasons.

For example, the scope of SAP adoption might be targeted at only a specific subset of business processes in the enterprise, while other parts of the business continue to function with minimal changes. In other cases, enterprise transformation is based on adopting both SAP and non-SAP packaged application solutions. For example, highly specialized, industry-specific packaged application solutions are adopted to gain a competitive edge.

Regardless of the specific scenario, in heterogeneous enterprises most SAP implementations must be integrated with a variety of non-SAP enterprise systems, portals, messaging infrastructure, security, systems of record, systems of engagement, and more.

The tooling included with SAP software addresses many of the requirements for creating SAP-specific environments. However, the classic approach to implementing even homogeneous SAP functionality generally leaves the business with a rigid solution that is difficult and expensive to change and enhance.

When SAP is used in a large, heterogeneous enterprise environment, SAP clients face a dilemma of selecting the correct set of tools and platforms to implement SAP functionality, and to integrate SAP with non-SAP systems.

The following questions are the most important to answer:

- What level of control of my data and processes do I want to maintain during implementation of SAP systems, and post-implementation?
- Will my enterprise approaches to mobile, business processes, integration, portals, security, management, monitoring, data, and business analytics work with SAP systems?

Organizations can enhance and extend pre-built capabilities in SAP software with best-in-class IBM enterprise software, to maximize return on investment (ROI) in their SAP investment, and to achieve a balanced enterprise architecture approach.

1.2 Critical success factors for an SAP-centric transformation

Large organizations around the world are pursuing enterprise transformations to streamline their business and IT landscapes by using prepackaged enterprise application suites, such as those provided by SAP. These transformations often replace a significant portion of the existing IT landscape, and it would not be possible to accomplish them without a prepackaged solution.

Packaged solutions often offer proven business functionality and processes. In addition, by replacing a difficult-to-manage IT landscape resulting from many decades of in-house systems development, acquisitions, and autonomous decisions, organizations are able to consolidate and simplify their processes and skills across the organization for better cost-efficiency.

The promise of replacing a complex IT environment with an *enterprise-in-the-box* solution that provides a pre-integrated set of configurable application modules, is appealing to business and IT leaders. However, as with anything that has the potential to provide great benefit, also several significant pitfalls that accompany large transformations based on packaged application solutions.

The following list includes some examples of such pitfalls:

- At some point after the initial implementation, the packaged applications need to be upgraded, either for new functionality or for continuity of support. If the packaged application has been heavily modified to meet the specific functional requirements of the organization, upgrade costs can be more than the initial implementation, making it a challenge to ever achieve a return on the investment.
- No packaged application provides a robust solution for every aspect of the business, and for every industry. As with all comprehensive solutions, there are many gaps in functionality, and many levels of maturity. One application module might rank a 10-out-of-10, but another module might rank only a 4-out-of-10.

  Even after a large packaged application transformation, most organizations continue to remain heterogeneous in nature. In addition, no business wants all of their processes to be based on a generic packaged process flow. Organizations must differentiate themselves to be competitive. Organizations need an additional platform to house their differentiating business processes and customizations, outside of the packaged solution, to avoid over-customization.
- These projects are often so large that they can eclipse all other discretionary spending within IT. When projects become that large, it is extremely difficult to balance architectural decisions that are in the long-term best interest of the enterprise with architectural decisions that specifically benefit only the transformation project.
- Often, an organization's data that is captured and stored by the packaged application is then subject to the licensing terms of that solution. If an organization is not careful in how it negotiates license terms, it can consequently lose control of its processes and data for future use and innovation.

Successful IT transformation investments in SAP systems establish an effective mechanism to ensure governance and balance throughout the lifecycle of the transformation, and beyond. This includes balance between short-term benefits and long-term benefits, and balance between what will benefit the project at the expense of the enterprise and what will benefit the enterprise at the expense of the project.

A more practical approach to SAP-centric enterprise transformations includes the following elements:

- Deploying a system of engagement to complement SAP transactional functionality as the system of record
- Balancing SAP with an application-independent, industry-leading integration platform
- Establishing governance for architecture decisions
- Avoiding custom coding

1.2.1 Deploying a system of engagement for SAP

Systems of engagement are systems built to connect users, mobile, the cloud, the web, partners, social media, and the Internet of Things (IoT), which includes billions of devices, to the systems of record that contain business data, business logic, business processes, and transactions.

Because SAP applications provide user interfaces (UIs) and some concept of *process*, most SAP implementations use the SAP transactional backbone (the *SAP system of record*) as their default system of engagement for the vast majority of their business processes.

This suboptimal approach marries the business processes of an organization and the associated business policies to a less-than-user-friendly platform that is bound by the constraints of information technology (IT) application lifecycle management (ALM). In most enterprises, this lifecycle typically spans three to six months, and in some cases a year or more. However, a wide range of business changes should optimally be implemented in a matter of days or weeks, not months.

Opportunities for continuous process improvement and business performance optimization are limited through the use of a system of engagement that is intrinsically bound by the IT application lifecycle.

Businesses today are increasingly demanding system of engagement agility, usability, and differentiation, not only in their external-facing applications, but in their core execution systems also. Externalizing at least some degree of SAP process control accelerates business-led change, enabled by a flexible process layer in the system of engagement. It can deliver dramatically enhanced flexibility, agility, and control over the traditional SAP implementation approach.

IBM Smarter Process helps to optimize both homogeneous and heterogeneous SAP processes by providing a best-in-class process orchestration layer external to SAP. These capabilities can be added to SAP at any time, both during the implementation and post-implementation.

1.2.2 Balancing SAP with an application-independent, industry-leading integration platform solution

Successful SAP-centric enterprise transformations adopt SAP systems as a set of business services that are integrated into the heterogeneous enterprise on service-oriented architecture (SOA) principles using best-in-class enterprise integration middleware. IBM calls this perspective a *system of record* (SOR), differentiated from the *systems of engagement* (SOE) that uses business services in the SOR.

Good selection of enterprise integration software is one of the most critical success factors for SAP adoption in a heterogeneous enterprise. The integration software must be able to support a diverse set of application environments, channels, and industry trends, not just SAP systems. The selection of the foundational software infrastructure must be based on middleware requirements for enterprise-class scalability, reliability, security, manageability, open standards, and application-independence.

SAP has been extending its offerings beyond the traditional scope of packaged enterprise application software into the area of integration middleware, including integration, business process management, mobile, portal, development lifecycle, business intelligence, database, and so on. These middleware offerings are typically well-integrated within the SAP business application suite, but are often not adequate for integrating SAP and non-SAP systems.

Organizations adopting SAP software should not assume that all of the middleware products that SAP offers are equally proven in the industry and are equally robust. Instead, a governance process for architecture decisions should be used to evaluate and choose enterprise-class software infrastructure to support all of the application environments.

For example, one of the misconceptions that can have significant negative effect on an enterprise's balanced middleware strategy in SAP-centric enterprise transformations is the belief that integrating external packages and components will "break" the SAP solution. This leads to the conclusion that the only viable option to integrate SAP software with an enterprise is to use SAP-provided integration middleware.

However, SAP has remained one of the most integration-enabled platforms for the last 20 years. It integrates well with robust third-party integration middleware software, such as the software provided by IBM.

Failure to establish a solid enterprise middleware strategy, and limiting the thought process to getting SAP into production, leaves optimization and enterprise alignment of the foundational software infrastructure out of scope. It also greatly reduces the chances of achieving long-term ROI from the SAP solution.

IBM proposes the following guidelines for implementation:

- First, use SAP as a set of business services in conjunction with process management technology, but with minimum customization.
- Next, adopt enterprise-grade integration architectures and technologies.

1.2.3 Establishing governance for architectural decisions

SAP adoption is a complex endeavor that involves multiple stakeholders, partners, vendors, technology alternatives, and so on. When a massive investment is directed toward a single vendor, such as SAP, organizations must establish an effective mechanism to make key architectural decisions in the best interest of their enterprise, and not just follow vendor bias.

Lack of visibility, and lack of well-documented architectural decisions with an enterprise perspective in mind, together represent a risk of filling the vacuum with project-level technical decisions that are often tactical in nature, and tend to represent an unbalanced opinion.

Establishing a governance mechanism for making balanced integration and process management technology decisions in the organization's own best-interest is important. It should prevent architectural decisions regarding application infrastructure from being driven by a vendor to push application-specific integration infrastructure into the enterprise domain. Good governance processes must ensure that carefully selected enterprise-class integration technology is used to integrate SAP with non-SAP systems.

Enterprise governance for architectural decisions must be established early in the SAP transformation program, and continuously used for probing and reviewing delivery teams to ensure correct technology usage.

1.2.4 Avoiding custom coding

The long-term ROI on SAP implementations is directly linked to the level of SAP customization. The premise of SAP application investments is *buy not build*. Organizations with significantly customized SAP deployments incur a much larger upgrade cost when compared to a more standard SAP installation.

Excessive levels of SAP customizations make the efforts required for version upgrades expensive, at times comparable with the effort of the initial SAP implementation. SAP customization that exceeds a certain level of ready-for-immediate-use business functionality might result in an unsustainable increase in future upgrade costs.

Excessive customization is characterized as *building a middleware within a middleware*. Excessive customization is the single biggest reason for not realizing long-term ROI on SAP investments, losing control over cost of ownership of the SAP solution, and not being upgradable in the future.

SAP implementation project leaders are often measured only on *getting SAP into production on time and within budget*. Usually, they are not measured on the ability to upgrade the resulting SAP implementation at a later date, which might not occur for five to seven years, and can be of a higher cost than the original implementation if the SAP solution was over-customized.

IBM Smarter Process technology provides application-independent technology to extend packaged applications without “breaking” the applications. IBM Smarter Process enables a balanced SAP adoption strategy: To consolidate and optimize non-differentiating, commodity-type business processes in SAP, while using a business process management (BPM) strategy to assemble and optimize differentiating business processes.

IBM Smarter Process technology also includes a robust business rules management system (BRMS) to produce new business logic that is not present in SAP systems, and not generally expressed as a business process. Complementing SAP solutions with BRMS enables organizations to reduce over-configuration and customization of the SAP environment, and to extend SAP and non-SAP functionality.

1.3 Combined value of IBM and SAP software

SAP is a market leader in packaged applications. IBM is a market leader in enterprise software. IBM is also a market leader in the delivery of SAP projects.

IBM enterprise software extends SAP value in multiple ways. The following sections explain the value of using IBM software in SAP solutions. They also explain how IBM software increases the value of an SAP-centric business transformation programs.

1.3.1 Reduced business and IT risk

IBM offers a powerful set of products that provides a wide set of enterprise functionality, and that is widely recognized as a best-in-class enterprise middleware platform. IBM middleware products provide mature, highly scalable integration technologies that have been proven globally across industries, customers, and geographies. IBM middleware has an extensive set of pre-built integration capabilities for SAP that have been extensively and successfully used in a large number of successful SAP projects.

The differentiating value of SAP enterprise middleware is pre-built, ready-for-use integration within SAP: SAP functional modules and bolt-ons, pre-built content for a portal, and mobile applications.

Successful SAP adopters reduce business and IT risk by combining two aspects of your enterprise infrastructure:

- The value of pre-built SAP integration within the SAP domain based on SAP middleware (*integration you buy*), also known as *inner ring*
- Best-in-class IBM enterprise middleware for custom integration in the enterprise that needs to be developed for an SAP-centric transformation program (*integration you build*), also known as *outer ring*

IBM uses the concepts of inner and outer ring to mean what happens inside SAP stays inside SAP, but what happens in the enterprise goes through IBM software.

1.3.2 Accelerated SAP integration into a heterogeneous enterprise

IBM software provides a powerful set of pre-built, enterprise-level capabilities that are pre-integrated with SAP systems, but are also integrated with non-SAP systems in the enterprise landscape. These include systems built on open standards, cloud-based offerings, non-SAP packaged applications, internally developed existing systems, and so on.

These capabilities enable IBM software to become a solid foundation for an open and heterogeneous enterprise architecture where SAP plays an important or, sometimes, the primary role. The solid foundation helps organizations to meet enterprise-level goals and imperatives, such as reuse of existing enterprise assets, use of standard SAP applications with minimum customization, loose coupling of SAP and non-SAP applications, use of open standards, and reuse of existing IBM infrastructure.

Examples of this extensive set of IBM software capabilities include areas such as mobile technologies, smarter process, decision management, enterprise content management (ECM), business analytics (BA), portal, master data management (MDM), ALM, security, systems management, and more.

Paring these IBM software capabilities with SAP solutions in large-scale enterprise transformations enables customers to accelerate SAP integration into the heterogeneous enterprise environments, and build an enterprise approach that is not restricted to SAP alone.

1.3.3 Business agility

SAP is rarely implemented with business agility in mind. Often the opposite takes place: SAP is used to harmonize, standardize, control, simplify, enforce conformance, and reduce costs. Business transformation based on packaged application software such as SAP is usually a huge project, disruptive to the enterprise, and does not differentiate from competitors running the same packages.

Business agility enables customers to implement frequent changes in business processes or business rules without changing the SAP solution. It also enables customers to implement business processes underpinned by SAP data, but which are only partially met by the SAP solution, or are not realizable in the SAP system without extensive customization coding.

Combining IBM Smarter Process, IBM MobileFirst, business analytics, and cloud (just as examples) with SAP solutions enables organizations to maintain business agility in SAP-centric transformations, and to maintain application innovations without “breaking” SAP.

1.3.4 Cost reduction

IBM software reduces the cost of SAP-centric enterprise transformation by using the correct best-in-class tools depending on the type of optimization needed.

For example, running a robust IBM application integration middleware platform with market-leading performance characteristics, can lead to substantial cost savings. These savings are realized from requiring less processor and random access memory (RAM) resources to handle enterprise-scale workloads when compared to less optimized middleware, which might require more hardware resources.

2

IBM Reference Architecture for SAP

Many organizations are both IBM customers and SAP customers in the global information technology (IT) industry. Because of this vendor overlap, organizations are repeatedly architecting IBM and SAP software products to work together in the best possible manner.

The purpose of this book is to capture many of the leading architectural practices when balancing a large SAP investment with IBM enterprise software technology. The book is meant to serve as a starting reference point that organizations can use to take advantage of IBM global presence, experience, and gathering of leading practices.

This chapter provides an overview of the IBM *reference architecture* for SAP, and describes high-level architecture patterns. The subsequent chapters in this book provide more details for several key technology domains that encompass the scope of this reference architecture.

This chapter includes the following topics:

- 2.1, “Overview” on page 10
- 2.2, “Architecture goals” on page 10
- 2.3, “IBM Reference Architecture for SAP overview” on page 12

© Copyright IBM Corp. 2014, 2015. All rights reserved.

2.1 Overview

A reference architecture is an asset that includes a collaborative set of architectural guidelines for use by all of the teams in an organization. Reference architectures provide a set of predefined architectures, also known as *patterns*, designed and proven for use in particular business and technical contexts. Reference architectures also include supporting artifacts to enable their use.

Having a documented reference architecture in place helps to drive alignment for the projects implemented across the enterprise. Without a reference architecture, each project team might decide to use the tools available to them in completely different ways, or they might implement their project using different methodologies or tools.

The purpose of IBM reference architecture for SAP (also referred to as the *reference architecture* in this book) is to inform and guide technical professionals, architects, and decision makers who are responsible for designing solutions that include IBM software that must coexist with SAP systems.

The scope of the reference architecture is the prescriptive use of IBM software in SAP implementation projects. The reference architecture is generic in that it does not depend on the specific type of SAP implementation. The reference architecture is focused on defining the use of IBM software with SAP solutions, and is not intended to address the internal aspects of the SAP components.

IBM Reference Architecture for SAP is based on the experience gained in IBM internal and client projects, with a focus on existing system modernization, large-scale SAP implementations, IBM internal transformation projects, and IBM strategies for IBM MobileFirst, API Economy, systems of engagement (SOE), business analytics (BA), IBM Smarter Process, big data, and cloud.

This reference architecture, although it is a prescriptive blueprint, can be implemented with multiple points of variability based on architectural decisions, timelines, budgets, skills, and other criteria. These points of variability do not change the prescriptive nature of the reference architecture. It is designed to maximize the value realized from IBM software products used in large-scale enterprise transformation programs based on an SAP solution. In particular, IBM software and SAP projects where an organization has a heterogeneous IT environment.

2.2 Architecture goals

The overall goal of the IBM reference architecture for SAP is to establish a prescriptive blueprint that can be used as a template when designing solutions that include IBM software and SAP systems. Following the guidelines provided by this reference architecture helps organizations to achieve critical success factors for SAP adoption in a heterogeneous enterprise.

2.2.1 Use standard, non-customized SAP applications

One of the primary business drivers for adopting a packaged enterprise resource planning (ERP) solution, such as SAP, is the business value achieved from using pre-built, ready-for-immediate-use components. SAP enables the application components to be tailored to fit a particular environment through configuration and customization.

Configuration changes can be made with minimal effort and in a version-safe manner. Custom development, however, can require more effort, and begins to take away from the value provided from a packaged application. Even modest customizations can lead to a chain of dependencies that inhibit taking advantage of new features and functions in SAP software, or upgrading SAP software cost-effectively in the future.

2.2.2 Reuse pre-built SAP integration

The overall architecture goal is driven by a need to use packaged application functionality (SAP) and packaged SAP integration (*integration you buy*) within the SAP domain while providing integration into the existing enterprise landscape outside of the SAP domain.

Even organizations with significant alignment with SAP rarely use SAP for more than 60% of their total business automation needs, according to IBM practitioners with extensive experience in SAP projects. This fact makes the enterprise integration of large-scale SAP implementations a key success factor of the overall SAP transformation program.

The goal of the reference architecture is to provide a blueprint for using market-leading IBM technology for large-scale SAP integration into the enterprise.

2.2.3 Use best-in-class technologies when extending beyond the SAP domain

Often, the best approach is to use SAP packaged applications and application infrastructure where SAP has a proven solution and it fits the needs of the business without significant customization.

However, in cases where data or processes need to extend outside of SAP, it is usually best to use industry-leading, application-independent software technology.

The IBM software portfolio is unique in the industry when it comes to providing a comprehensive selection of middleware platform technology to complement SAP systems in heterogeneous environments. IBM software provides organizations with consistency, scalability, reliability, flexibility, and asset reuse for the application software infrastructure needs across all application domains throughout the enterprise.

2.2.4 Use open, well-established standards

Standards-based implementation of integration logic greatly reduces the chance of vendor lock-in. It also increases the availability of tools that easily connect to the solution, and that provide support for development, monitoring, and testing.

Various well-established standards exist across industries at the domain, message, and protocols layers. Proprietary middleware platforms should be avoided to retain flexibility and ownership of any custom-built logic or functionality.

2.2.5 Use pre-built software capabilities provided by IBM

The IBM Reference Architecture for SAP also provides a blueprint for using enterprise capabilities of IBM software that complement the SAP platform to provide a complete enterprise transformation solution.

In many cases, IBM software provides enterprise capabilities that are more robust than the corresponding capabilities in SAP software, for example, business process management

(BPM), IBM MobileFirst, enterprise integration, IBM DevOps solution, IBM Enterprise Content Management portfolio (IBM ECM), and so on.

In other cases, IBM software enables organizations to enhance and extend existing SAP capabilities into a heterogeneous enterprise environment, for example business analytics and cognitive computing.

In yet other cases, IBM software capabilities do not exist in SAP systems. Introducing IBM software in the solution enables organizations to differentiate the SAP solution from similar SAP implementations in other organizations. An example of such a unique IBM software capability is adding business agility to the SAP solution with IBM Smarter Process.

For each such capability of IBM software, the reference architecture provides a separate architecture component that describes how the IBM software capability should be integrated with the SAP system to provide a complete enterprise transformation solution. A key point is that IBM software excels when adding SAP systems to a heterogeneous enterprise as another service provider. IBM Smarter Process for SAP adds value equally to both homogeneous and heterogeneous SAP processes.

2.3 IBM Reference Architecture for SAP overview

IBM Reference Architecture for SAP is defined through various levels of abstraction, or *views*, thereby providing more flexibility in how it can be used.

2.3.1 Systems of engagement, record, and interaction

From an architecture perspective, the IBM strategy revolves around three types of *systems* (engagement, record, and interaction), as shown in Figure 2-1.

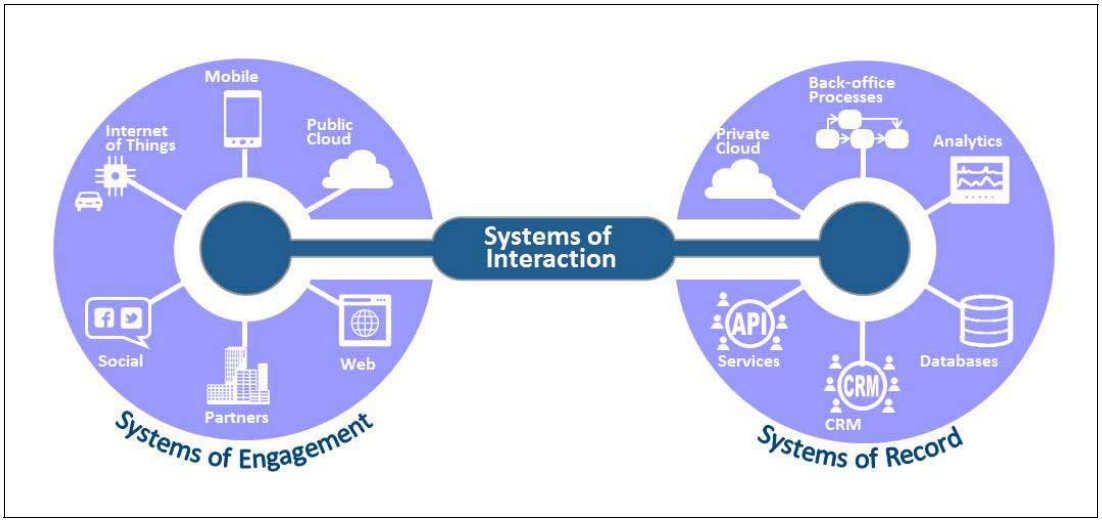

Figure 2-1 Systems of engagement, interaction, and record: Logical view

Systems of engagement (SOE) are systems built to connect to users, mobile apps, the cloud, the web, partners, social media, and the Internet of Things (IoT), which has billions of devices. SOE will continue to become more diverse, and drive new business value to agile enterprises.

An example of new business value is application programming interface (API) management, where organizations publish APIs, such as Uniform Resource Locator (URLs), that are consumed by application developers building apps. These apps use the APIs in non-traditional usage patterns, billing models, and application types.

SAP, existing, database, and other third-party systems constitute transaction-oriented *systems of record*. They are traditionally designed around discrete pieces of information or records. The benefit of systems of record is that they provide business data, business logic, business processes, and transactions to support day-to-day business operations.

Systems of interaction provide connectivity between SOEs and systems of record.

Figure 2-2 shows the architectural view of systems of engagement, interaction, and record.

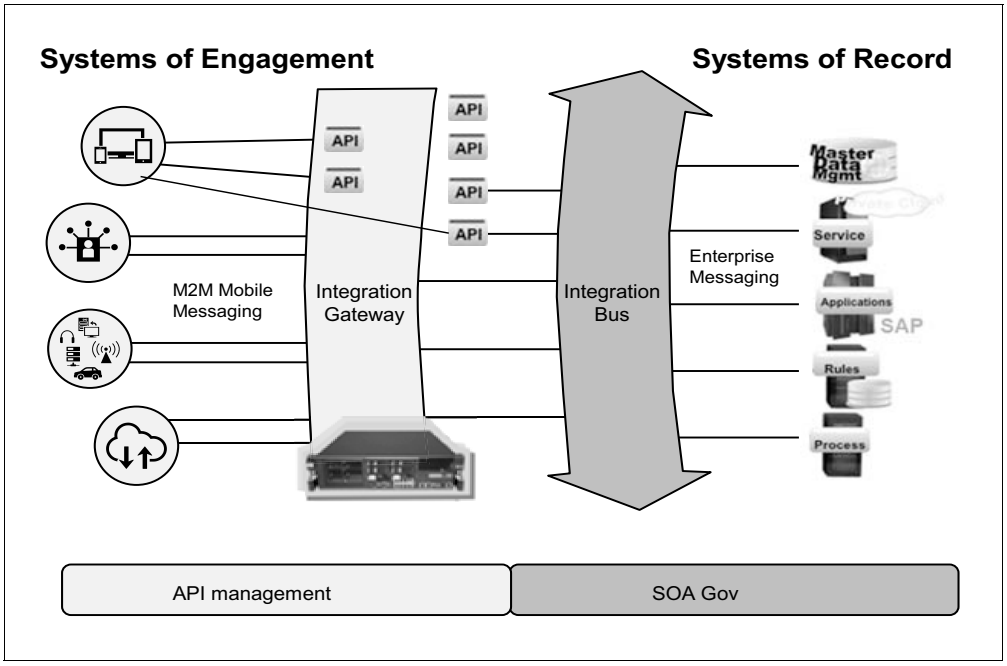

Figure 2-2 Systems of engagement, record, and interaction: Architecture view

The architecture view shows the need for the heterogeneous enterprise to focus on systems of engagement beyond a single technology approach tied to a system of record.

Integration gateways are a critical part of the architecture. They provide users registration, security, message validation, API management, and routing. Additionally, an integration bus bridges the integration gateway and the systems of record.

An important note is that governance, reviews, and management need to be considered across the project and technology lifecycles. This is another reason to carefully separate adopting SAP as another system of record from an enterprise solution to support systems of engagement.

2.3.2 Services view

This section provides a services view of IBM Reference Architecture for SAP, and describes key components in the SAP-centric heterogeneous enterprise architecture.

Figure 2-3 shows services of the reference architecture. Note that the logical view shown in the figure represents a superset of the functional components that might be involved in any given instantiation of this reference architecture in a given project.

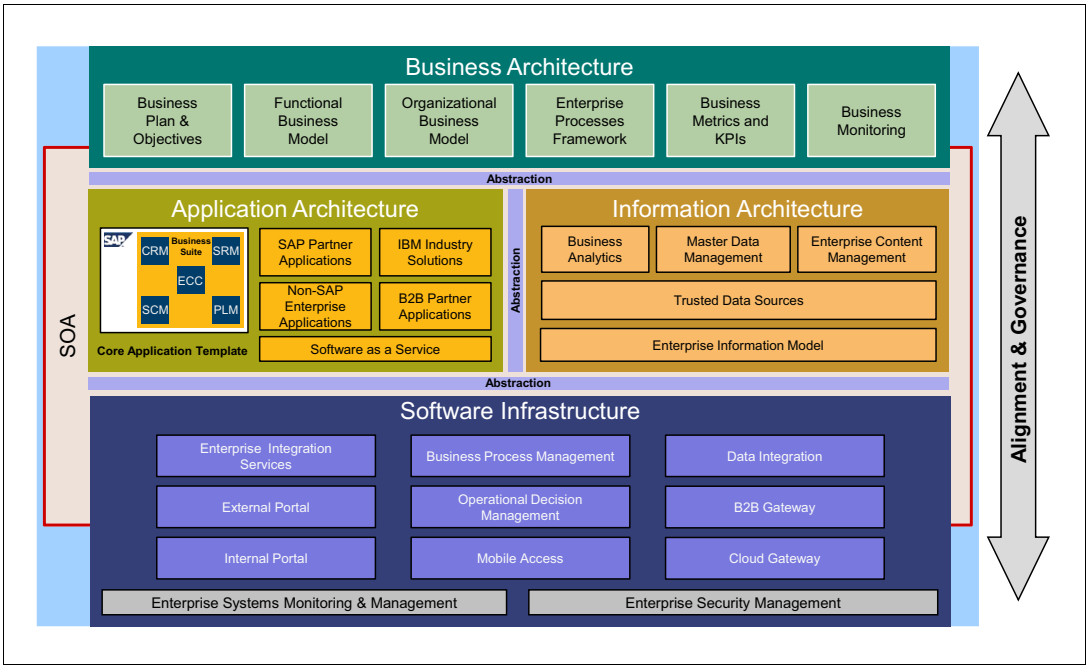

Figure 2-3 IBM Reference Architecture for SAP: Services view

Business architecture

Business architecture is the primary link between business and IT. Business architecture defines artifacts, standards, and principles that are used as a trusted source by executives, architects, and developers to align enterprise initiatives and solution content with business strategy.

This reference architecture does not completely define the business architecture. However, it outlines a core set of business architecture services that are normally expected to be already established in an enterprise that embarks on a large-scale SAP-centric transformation.

Most organizations have well-defined organizational and functional models. However, more complex organizations are adopting more matrix business architecture models, and establishing an *enterprise process framework* (EPF) that defines a trusted and clear taxonomy and ownership for the overall management of the enterprise business.

The scope of the enterprise transformation program is often defined within the EPF, and provides a foundation for governance in a complex transformation program. The governance model provides clear direction, focus, and executive commitment. Because the program is enterprise-wide and can be global in scope, the program structure based on an EPF enables better engagement of senior executives and regional leadership.

Information architecture

The information architecture provides guidance, templates and reusable artifacts, and information services from which to build information deliverables for a specific enterprise solution.

The information architecture is composed of an enterprise information model, trusted data sources, and enterprise information services, including master data management (MDM), Enterprise Content Management (ECM), and business analytics (BA). It defines the business objects, processes, and data, and establishes their inter-relationships at different abstraction levels.

The *enterprise information model* is an enterprise-wide, platform-independent, and conceptual data model for the data terminology of the business. This model defines a consistent terminology, and depicts, using entity relationship diagrams (ERDs), the business rules that govern the relationships between the business terms in the business process areas. The enterprise information model should be used as a base for creating conceptual data models for specific solutions.

An SAP data model does not supersede an established enterprise information model.

Trusted data sources define the optimal source of data for specific subject areas and business areas of data. The identified strategic sources should be the first choice for initiative or solution data repository use. Typically, five types of trusted data sources exist:

- Master data stores
- Transactional data stores
- Operational data stores
- Data warehouse
- Data marts

Applications architecture

The heterogeneous enterprise includes SAP and different types of non-SAP applications. The following list includes different dimensions of non-SAP applications. These dimensions are complementary, and are described here because they drive different SAP integration considerations:

- SAP application modules. Packaged applications provided by SAP.
- SAP partner applications. Non-SAP vendor applications, certified and recommended by SAP, which fill functional gaps or implement functions in the SAP portfolio.
- IBM Industry Solutions. Application solutions provided by IBM that are complementary to SAP, for example, IBM Commerce.
- Non-SAP enterprise applications. Enterprise portfolio of applications that are not included in the scope of the SAP transformation program (they are not replaced by SAP software).

  This category of application might require refactoring, because part of their original functionality might have to be moved to SAP systems. Such refactoring is different from mere SAP integration at the technical interface level, because functional changes inside non-SAP applications can require a significant effort.
- Software as a service (SaaS). Represents cloud-based business services.
- Business-to-business (B2B) partner applications. External applications connected by B2B gateway.

Services scope in SAP adoption project

During an instantiation of this reference architecture, the services in scope are categorized based on the expected effect on the SAP adoption project. This categorization enables distinguishing the components under direct control of an SAP adoption project from those under enterprise architecture control.

This categorization has important implications for the overall success of the SAP-centric enterprise transformation. It includes all of the components of the solution. Therefore, it enables your organization to provide the proper attention and planning to transform critical non-SAP components in the heterogeneous enterprise that are essential to support a new SAP solution.

The services are categorized in the following way:

- The first category includes *services that are directly implemented by an SAP adoption* project. For example, enterprise integration services and the corresponding SAP integration patterns normally are developed within the scope of an SAP adoption project, even though these are non-SAP components. IBM Smarter Process for SAP adoption ideally starts at or before the SAP business blueprinting, and continues throughout the SAP implementation project.
- The second category of components already exists in the enterprise before an SAP adoption project, and are not affected by the SAP adoption. The *SAP implementation will be integrated with such components, but it will not change them*. For example, enterprise single sign-on (SSO) or enterprise directory, which are part of enterprise security infrastructure, are not normally in the scope of the SAP adoption project.
- The third category of components is an intermediate category which is not directly in the scope of the SAP adoption project, but *might require a significant transformation as a result of the effect that the SAP adoption has on the enterprise IT landscape*.

  For example, MDM might require enhancement or even partial redesign to incorporate SAP data into the enterprise MDM solution. In some cases, an MDM refactoring project might need to precede large-scale SAP adoption, or run in parallel, but in either case it must receive adequate focus in the enterprise.

2.3.3 Application integration: Inner ring and outer ring architecture

This perspective of the reference architecture is focused on a federated approach to integration implementation. It includes the same application, information, and software Infrastructure services as in the logical view of the reference architecture.

IBM Reference Architecture for SAP is based on two federated technology domains (also known as *rings*), as shown in Figure 2-4 on page 17.

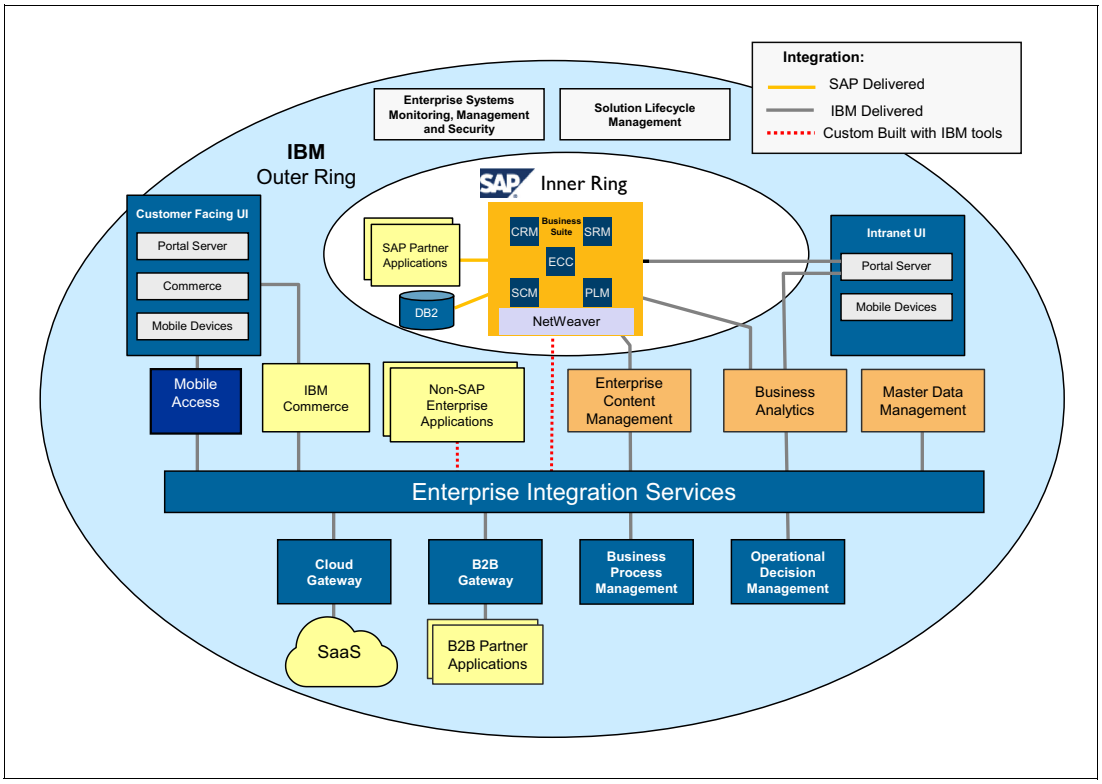

Figure 2-4 Inner ring/outer ring architecture

Inner ring

The inner ring is the SAP technology domain and includes applications, technology, and integration purchased from SAP and SAP business partners. This reference architecture assumes that it is not practical to force SAP consultants to use non-SAP technology for packaged content or customization within the SAP inner ring, even if it is technically feasible.

Therefore, the reference architecture is designed to encourage the use of SAP middleware tools and technology within the inner ring, but use robust application-independent technology whenever data or processes move into the outer ring. In this way, SAP customers can achieve optimal efficiency, maximum flexibility, high reliability, and the least amount of risk during large-scale transformation projects.

Outer ring

The outer ring represents the entire technology domain outside the cluster of SAP applications. The outer ring includes existing applications, packaged applications from software vendors other than SAP, non-SAP applications hosted in the cloud, and all other shared software infrastructure platform technology.

Previously implemented installations of SAP, or acquisitions that might still exist within the enterprise, are a gray area. Normally, these older versions of SAP, such as SAP R/3 instances, are classified as existing systems, and categorized into the outer ring domain.

The key to the outer ring is the reuse of common, enterprise-class, application-independent, software infrastructure across the heterogeneous IT landscape. IBM provides the best channel to building this layer of consistency, resiliency, and flexibility of software infrastructure, rather than acquiring the various software components from different vendors and dealing with incompatibility issues.

The following list includes other key characteristics of outer ring software infrastructure components:

- Functionally rich. The software infrastructure component must be shared across all of the application environments for current and future products. For this reason, it must support a wide array of features and functions so that it can handle at least 90% of all enterprise requirements for the particular technology domain that it fulfills.
- Compatible with SAP applications and infrastructure. This reference architecture assumes that the organizations adopting it will be making a substantial investment in SAP applications. For this reason, the software infrastructure component should be designed to interoperate with SAP applications and software infrastructure within the inner ring domain.
- Reliable. The software infrastructure components must be highly reliable. For example, they should be able to run in a redundant, active-active configuration for months at a time without suffering crashes, needing to restart, or experiencing other types of outages.
- High-performance. It must make cost-efficient use of the hardware it is deployed onto. For example, if a platform requires 5x more hardware to perform similar tasks as on alternative technology, it should not be considered as an option for an outer ring software infrastructure component.
- Manageable. It should have dashboard-style management capabilities, so that multiple instances of the technology can be centrally managed from any location. In addition, it should support third-party system management consoles.
- Flexible. It should maximize the use of open standards where possible to enable cross-component interoperability and vendor independence. For example, process-oriented technology should support Business Process Modeling Notation (BPMN), service interfaces should support Web Services Description Language (WSDL), identity management should support Lightweight Directory Access Protocol (LDAP), and so on.
- Industry-leading. The software infrastructure components in the outer ring should be ranked by mainstream industry analysts within the top providers of application-independent technology for its class. Rather than spending months doing a thorough analysis of feature and function comparisons, validating that the software components under evaluation are among the top industry leaders is a good sign that the technology has enterprise-class characteristics.
- Longevity. The vendor of the software infrastructure component should have a long-standing history in the market as a leader in the particular domain of software infrastructure, such as enterprise integration, BPM, portal, mobile, MDM, ECM, and so on. Switching enterprise platform strategies is too expensive and disruptive to trust these decisions to small, up-and-coming vendors.

If an organization enables each project to make their own vendor and product decisions regarding the software infrastructure technology, pretty soon there will be hundreds of vendor products and methodologies in place across the enterprise. Over time, this approach results in extreme cost inefficiency and inflexibility. The inner ring/outer ring architecture ensures consistency and reuse across all future projects in the enterprise, SAP and non-SAP. The architecture balances simplicity with flexibility and control.

Each particular technology domain of the inner ring/outer ring architecture is addressed as a separate reference architecture in each chapter of this book.

2.3.4 Enterprise integration services

The scope of integration in this reference architecture includes the following domains:

- Initial data load into SAP systems
- Ongoing application and data integration of SAP and non-SAP systems

Initial data load into SAP systems

SAP implementations often require migrating large volumes and varieties of data from a range of systems that span multiple business functions. Poor-quality data often hinders these SAP initiatives, causing project delays, cost overruns, and ultimately poor adoption of SAP systems across the organization.

IBM InfoSphere Information Server Ready to Launch for SAP Applications provides a complete, scalable, repeatable, and reusable SAP-specific data migration solution. Built on IBM InfoSphere Information Server technology, it consists of a comprehensive, documented data migration methodology that is specific to SAP migrations and consolidations, plus a set of data migration accelerators.

The following list describes some of the benefits of IBM InfoSphere Information Server Ready to Launch for SAP Applications:

- Provides a complete data migration solution for SAP applications that couples data migration leading practices with software and services from a leading global SAP implementer
- Helps reduce the risk and time of SAP migrations and consolidations by taking data off the critical path and establishing standardized, repeatable, and data-centric processes
- Turns SAP data migration expense into a strategic investment by delivering a repeatable and reusable infrastructure that can be used for multiple SAP rollouts
- Improves SAP business process execution and customer experience by increasing information accuracy

This solution is well-documented, and encapsulates a breadth of SAP experience with proven methodology for data migrations and consolidations used in thousands of successful client SAP implementations. These SAP implementations were led by IBM Services, the largest worldwide systems integrator, business partner, and software integration partner for SAP.

Ongoing data integration with SAP

The ongoing data synchronization is the next major data integration activity in the SAP enterprise transformation, after the one-time, initial data load. This activity involves all of the interfaces between existing systems and the new system. It is run on a continuous basis after the new system is live.

The ongoing interactions encompass a wide range of integration requirements and approaches: Batch versus transactional, synchronous versus asynchronous, service-oriented versus message-oriented, and so on.

In the SAP implementation methodology, ongoing data synchronization is captured as interface definitions between SAP and non-SAP systems. Such interface definitions are identified during the blueprint phase of the SAP methodology.

It does not differentiate based on the technical nature of the interactions, for example, transactional versus batch. Alternatively, IBM software provides different technologies for addressing the different integration requirements, and the available technologies provide partially overlapping capabilities.

To provide better alignment with the SAP implementation methodology, this reference architecture treats both transactional and batch integration aspects of ongoing SAP integration with non-SAP systems as a single architecture component. Enterprise integration services provide a unified set of integration patterns, as described in Chapter 3, “Enterprise integration services for SAP” on page 39.

The reference architecture for the enterprise integration services components is shown in Figure 2-5.

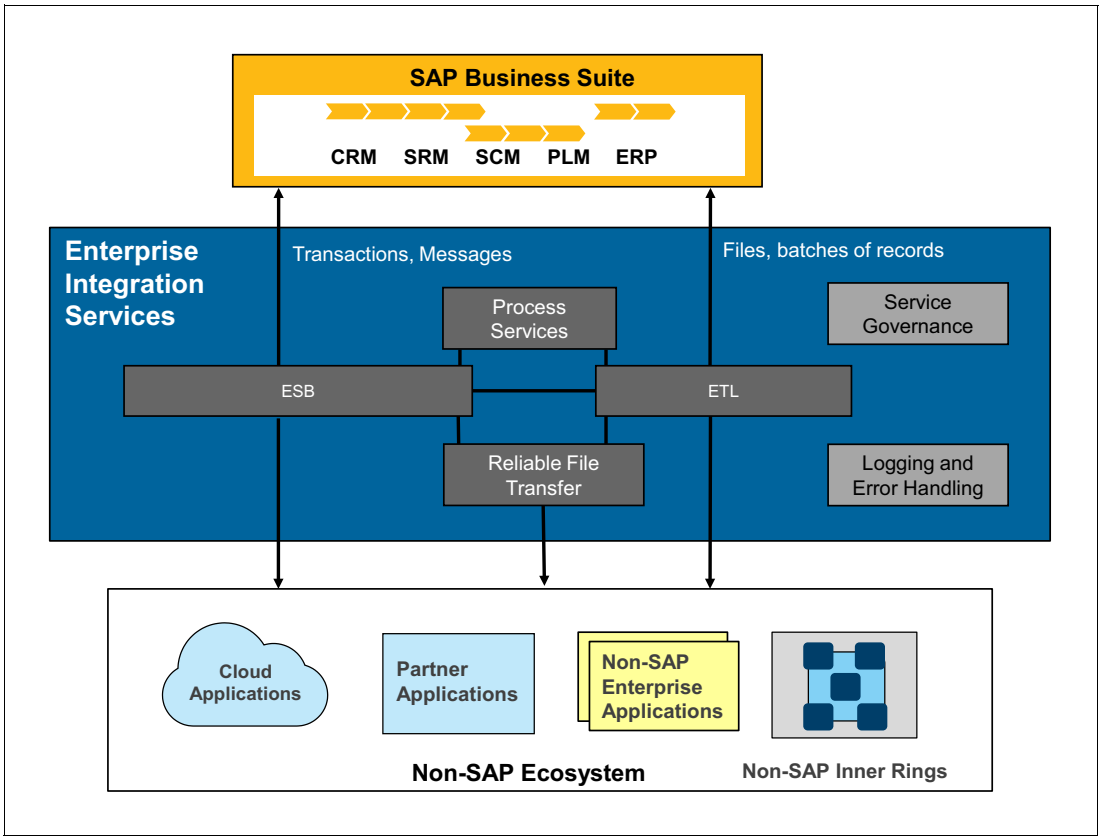

Figure 2-5 Enterprise integration reference architecture

Several components make up the enterprise integration services architecture:

- Enterprise service bus (ESB)
- Extract, transform, and load (ETL)
- Service governance
- Reliable File Transfer (RFT)
- Process services
- Logging and error handling

These components work together to provide the capabilities required to connect SAP with non-SAP applications within the enterprise, business partners, and cloud-based applications.

The ESB component is responsible for providing *connectivity and integration logic for transactional interfaces*. The primary function of the ESB is to decouple and isolate the application endpoints from one another, increasing the flexibility of the system and reducing the overall cost of integration.

ETL is a term used broadly to refer to the activities required to move large volumes of data between systems in large batches. In the context of enterprise integration services, the ETL component is responsible for providing *connectivity and integration logic for batch-oriented interfaces* for ongoing integration (as opposed to initial data load or conversion activities). As the name suggests, three major processes are involved in an ETL flow:

- Extract the data from the source system.
- Transform (and optionally cleanse) the data.
- Load the resulting data into the destination system.

ETL technologies are built to efficiently process very large sets of data, with internal staging of the data and parallel processing.

Service governance includes two major aspects:

- Service lifecycle management
- Service run time

In practical terms, the service governance component provides a *service repository* for storing service artifacts and a *customizable, ready-to-use process* for managing those service artifacts throughout the project lifecycle. In addition, service governance includes a *runtime service registry*, where the defined service policies and configurations are enforced by the service integration components, and the resulting runtime information is provided back to report on the service use.

RFT technology provides *central configuration and set up, centralized logging and monitoring* of all file transfers, and a standard solution with *established quality of service* (QoS) characteristics for implementing file transfers within the enterprise.

Process services are technical integration processes that provide *advanced integration logic* beyond the typical mediation that is provided by ESB components. Process services are typically implemented on a BPM platform.

A clear delineation should be drawn between business processes that provide the logic for business operations (including human interaction by business users) and process services that *satisfy technical integration requirements*. Process services can involve human interaction in some cases, but the users involved in those activities are typically in technical support roles.

For more information, see Chapter 3, “Enterprise integration services for SAP” on page 39.

2.3.5 Process optimization

With the mainstream adoption of IBM Business Process Management and IBM Operational Decision Management solutions, SAP-centric organizations now have a myriad of approaches to deploy and optimize business processes.

Organizations can choose any of the following approaches:

- Choose to deploy the rather loose concept of process imposed by the classic SAP configure/customize/document approach.
- Apply a light layer of workflow management using a tool such as SAP Business Workflow to complement the classic SAP configure/customize approach.
- Apply a comprehensive layer of BPM and ODM capabilities to provide deep business optimization potential for SAP, non-SAP, and heterogeneous SAP processes.

Migration to the SAP-packaged application platforms enables organizations to buy, configure, and deploy pre-built functionality rather than modify, enhance, and maintain existing applications. Adoption of the new business processes mandated by SAP fixed or configured functionality enables an organization to withdraw from service large numbers of inflexible existing application systems and replace them with a better alternative, albeit with trade-offs.

Because of the large single scope and complexity of most SAP implementations, however, SAP adoption introduces a new set of operational and organizational challenges, risks, and pitfalls.

Historically, BPM and ODM have represented an entirely different approach to business process optimization, an approach that is not disruptive and is designed to achieve business differentiation across a relatively small number of business processes.

This approach enables organization to selectively optimize any business process, addressing areas of concern across processes in priority sequence with small and incremental projects. BPM and ODM are designed to deliver complete visibility to process performance across core and fringe application systems.

With the advent of IBM Smarter Process for SAP, however, the historical positioning of BPM and ODM as an optimization tool for only a relatively small number of mission-critical or highly painful heterogeneous business processes has changed dramatically.

External BPM orchestration of a large number of homogeneous SAP processes is now feasible, thanks to an innovative IBM technology that enables business process designers to design, build, and deploy powerful external orchestrations of SAP processes without IT development or coding.

Rapid, flexible, business-friendly integration with the SAP graphical user interface (GUI) and technical integration engines now removes most of the historical barriers to large-scale BPM adoption in SAP environments. This integration delivers enormous business optimization potential, both within SAP and across the functional landscape.

The same IBM Business Process Manager technical capabilities that enable a broad-scale process orchestration of SAP and heterogeneous processes also help to reduce the amount of SAP customization and configuration. Avoiding customizing SAP solutions as much as possible is one of the key success factors of SAP-centric enterprise transformations.

A significantly customized SAP deployment incurs a much bigger upgrade cost compared to a more standard SAP installation. Excessive levels of SAP customizations can make the efforts required for version upgrades prohibitively expensive, sometimes on par with or even exceeding the effort for the initial SAP implementation.

Among the key reasons for customizing SAP are a business need to differentiate from competitors who also implement SAP solutions, and to provide functionality unique to a particular offering, market segment, or customer. IBM Business Process Manager and IBM Operational Decision Manager can be used to run, complement, and extend packaged SAP application processes to deliver a comprehensive business optimization platform.

At the same time, they can help realize important business requirements without over-customizing SAP. Figure 2-6 on page 23 shows how IBM Business Process Manager and IBM Operational Decision Manager optimize, complement, and extend SAP through a set of prebuilt and pre-integrated functional components in the IBM Business Process Manager and IBM Operational Decision Manager suite. These components enable a business differentiation for SAP adopters while avoiding excessive SAP customization.

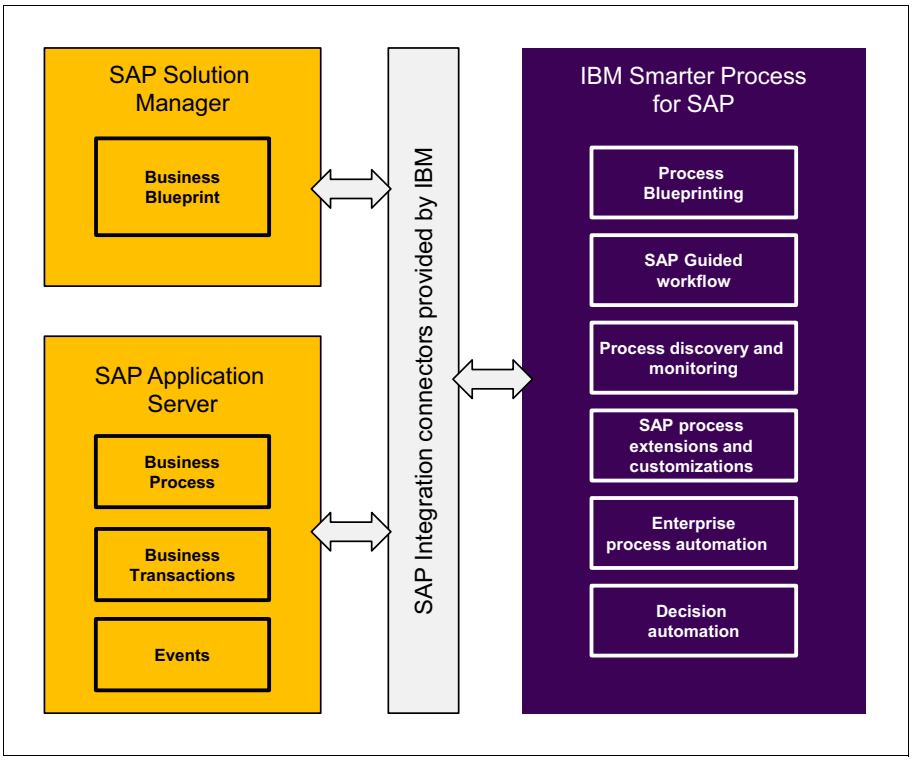

Figure 2-6 IBM BPM and IBM ODM extend and complement SAP for business differentiation

SAP business blueprinting in IBM Business Process Manager reduces SAP blueprinting time, cost, and risk by using an iterative, experiential-based approach to accelerate traditional SAP blueprinting.

Understanding packaged SAP processes and gaps can be difficult with only SAP tooling. Process blueprinting in IBM Business Process Manager enables you to import business blueprints from SAP Solution Manager. Then you can understand, edit, develop, configure, and customize the SAP business process blueprint. It uses state-of-the-art graphical modeling tools.

Then, you can export the blueprint back to the SAP environment. IBM provides pre-built integration between SAP Solution Manager and IBM Business Process Manager modeling tools to deliver this automated model exchange.

The process transaction flow documented in SAP Solution Manager does not necessarily guarantee that it is how a process actually works inside SAP. SAP Solution Manager documents the *expected* order of SAP transactions that users will start to support a business process. However, it is usually possible that users might start transactions in a different and unexpected order.

IBM Business Process Manager Guided Workflow for SAP has the capability to wrap a set of SAP transactions (native SAP screens) that constitute part of an SAP process with an automatically generated workflow. It guides SAP users through the correct sequence of SAP transactions for each process instance, while gaining real-time insight into business performance issues and opportunities.

Process steps that complement SAP functionality, such as approvals, escalations, exceptions, value management, and activities in non-SAP applications can be easily added to the automatically generated IBM Business Process Manager Guided Workflow for SAP steps. This capability enables an organization new to SAP to dramatically reduce deployment time and improve the adoption curve.

Experienced SAP organizations will find that guided workflows can quickly and easily improve the productivity, visibility, agility, and consistency of the SAP process execution. IBM Business Process Manager also helps to ensure that the end-to-end business process complies with service level agreements (SLAs), regardless of whether process steps are implemented in SAP.

The IBM Business Process Manager approach is the most effective option to help organizations avoid the number one pitfall when implementing packaged applications: Over-customization of the packaged content. Excessive customization can suppress the ability to perform functional application upgrades at a reasonable cost, or the ability to efficiently add new SAP packaged application modules and components in the future.

Process discovery capabilities are used to mine SAP business events to help discover how SAP processes are actually run, based on the sequence of business events that these processes generate. Process monitoring capabilities provide continuous business activity monitoring for any SAP process instance, and enable timely responses to business challenges.

IBM provides pre-built integration capabilities that enable organizations to complement SAP systems with business activity monitoring. These pre-built integration capabilities are non-intrusive, and are reusable for both SAP and non-SAP components.

Decision automation implemented by IBM Operational Decision Manager complements, and sometimes replaces, SAP business logic to improve business performance, deliver more agile business processes, introduce dynamic business functionality, and enable business-led change.

IBM Operational Decision Manager treats business rules as an enterprise-level asset, not a project level artifact, and therefore facilitates reuse across applications, organizational units, and business processes. IBM Operational Decision Manager business rules are not tightly coupled with any particular back-end application or technology, such as SAP systems. Therefore, they are equally well-suited for SAP and non-SAP process steps.

An enterprise-wide implementation of IBM Operational Decision Manager creates a single point of management for business rules used in both SAP and non-SAP enterprise applications. IBM Operational Decision Manager is based on open standards, providing flexibility and reuse of existing enterprise assets. IBM Operational Decision Manager seamlessly integrates with the native SAP business rules framework.

For more information, see Chapter 4, “Process optimization for SAP” on page 75.

2.3.6 User interface

Two domains of user interface (UI) types are needed in SAP transformation programs:

- External facing UI (customers, partners)
- Internal facing UI (employees)

The UI component in the IBM Reference Architecture wrap underlying SAP functionality with open standards, and enable you to include non-SAP UI content.

Multiple user interfaces to SAP are exposed through different UI channels, depending on the functionality required (see Figure 2-7).

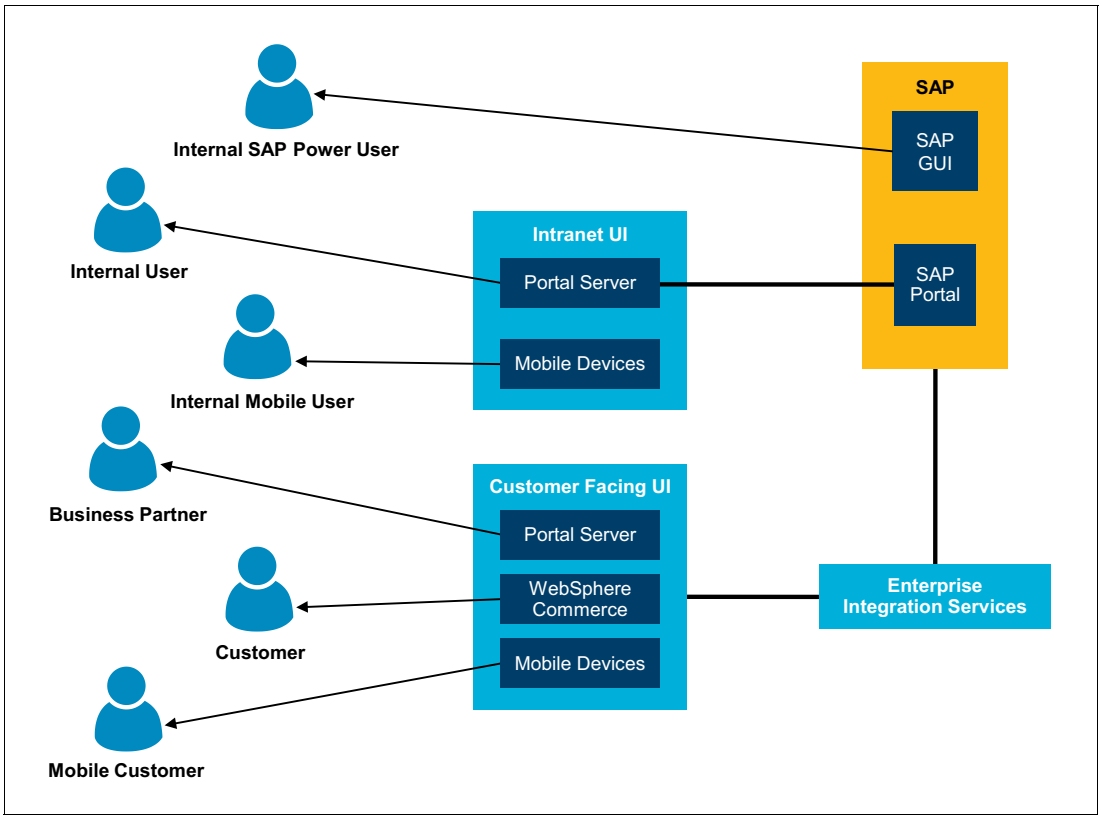

Figure 2-7 User interface channels in SAP-centric transformations

Internal SAP users can be divided into two classes:

- SAP power users. They perform their primary work in the SAP environment. These users have deep skills as SAP users, and typically are back-office workers. Power users use the native SAP GUI or an intranet portal. SAP power users typically represent a minority of the enterprise population.
- SAP occasional users. They require only on-and-off interaction with the SAP system, for example, a sales person who needs to check an order status. Occasional SAP users do not have deep SAP skills. They access SAP functions through an intranet portal, which wraps content from the SAP Portal (users do not access the SAP Portal directly). Occasional SAP users typically are the majority in the enterprise, even for enterprises with a high level of SAP adoption.

Customers and business partners use external portals that wrap SAP functions in open standards, and provide a company-standard and branding (non-SAP UI) “look and feel.”

Online commerce is a sophisticated and dynamic discipline. It is no longer simply about selling online; it is about delivering a consistent shopping experience across all customer touch points, including mobile, social, and in-store.

Separating the online experience from the back-end SAP processing is key to enable the overall solution architecture to take advantage of the specialized commerce software. IBM Commerce is a market-leading online commerce software platform. It enables organizations to evolve concurrently with ever-changing user experience expectations, emerging web marketing capabilities, and other trends.

IBM Commerce enables organizations to deliver a seamless, omni-channel shopping experience through contextually relevant content, marketing, and promotions, while extending their brand across digital and physical channels.

Mobile

Both internal and external users often require mobile access to SAP functions. Mobile access to SAP systems typically has a different set of requirements for internal and external users. The mobile UI for internal users typically provides a set of specific business functions, for example, labor claims. Requirements for the mobile UI for external users are typically more sophisticated, and typically require a differentiating user experience.

A key architectural decision in this reference architecture is to reuse an enterprise mobile platform for SAP mobile development, based on the principle of separating the enterprise mobile platform from vendor-specific server runtime environments. The IBM MobileFirst Platform is a best-in-class enterprise mobility platform built on open standards and designed for heterogeneous environments, both SAP and non-SAP back-ends.

The SAP mobile offering, SAP Mobile Platform (SMP), provides differentiating value through a set of pre-built mobile applications for SAP. SMP should be considered in a heterogeneous enterprise only when pre-built SAP applications can meet business requirements *as is*, with changes enabled through only SAP-supported configuration options. The SMP run time should not be used for custom development in a heterogeneous enterprise. Instead, it should be considered as a *black box* to support purchased mobile content from SAP.

Figure 2-8 shows a reference mobile solution for SAP systems that consists of two technology domains:

- The standard enterprise mobile platform domain for all SAP and non-SAP enterprise mobile applications based on IBM MobileFirst
- The *black-box* SAP mobile domain used exclusively for deploying pre-built SAP mobile applications that only meet business requirements *as is* or through supported configuration

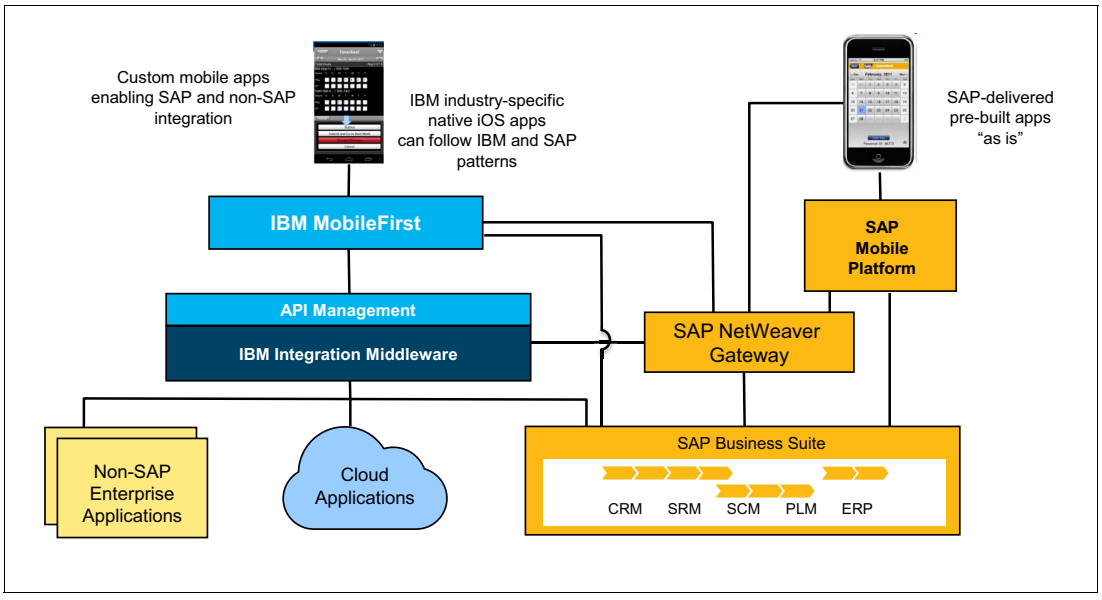

Figure 2-8 Reference mobile solution for SAP

IBM MobileFirst is a market-leading enterprise mobility platform built on open standards and designed for heterogeneous environments, both SAP and non-SAP back ends.

Using IBM MobileFirst does not merely enable access to SAP data through SAP integration capabilities that are provided by IBM. The IBM MobileFirst portfolio provides a comprehensive set of capabilities needed for enterprise mobile enablement, such as application development, device and application management, mobile analytics, mobile security, IBM DevOps solution integration, and many others.

IBM MobileFirst is designed to support any type of systems of record, rather than favoring one system or technology in particular, such as SAP systems. When gaps exist between SAP-provided mobile applications and business requirements that cannot be addressed through mere SAP configuration, custom development on the SAP Mobile Platform (SMP) should be avoided in a heterogeneous enterprise. Excessive SAP customizations result in significant problems with future platform upgrades.

Custom development of mobile applications for SAP is fully enabled with the IBM MobileFirst Platform, which should be the standard enterprise platform for any custom mobile development. IBM MobileFirst provides pre-built integration with the SAP NetWeaver Gateway, and a rich set of SAP integration options made available by re-using IBM integration middleware.

Project experiences show that the cost of custom development for SAP on IBM MobileFirst is significantly lower compared with other types of custom application development, for example, full-featured web-based business applications. IBM MobileFirst not only provides a best-in-class mobile development platform, but it includes a rich set of fully integrated enterprise capabilities for security, lifecycle management, mobile analytics, user experience feedback, and many others.

For more information, see Chapter 5, “Mobile access for SAP” on page 117.

Portal

IBM WebSphere Portal is the market-leading platform for delivering relevant, personal, and engaging user experiences to customers, partners, and employees. By integrating best-in-class business applications from SAP, with leading digital experiences from IBM, organizations can compete more effectively and enhance the productivity of their employees.

Figure 2-9 shows how IBM WebSphere Portal can effectively integrate with SAP applications, SAP Enterprise Portal, and with non-SAP enterprise and cloud applications.

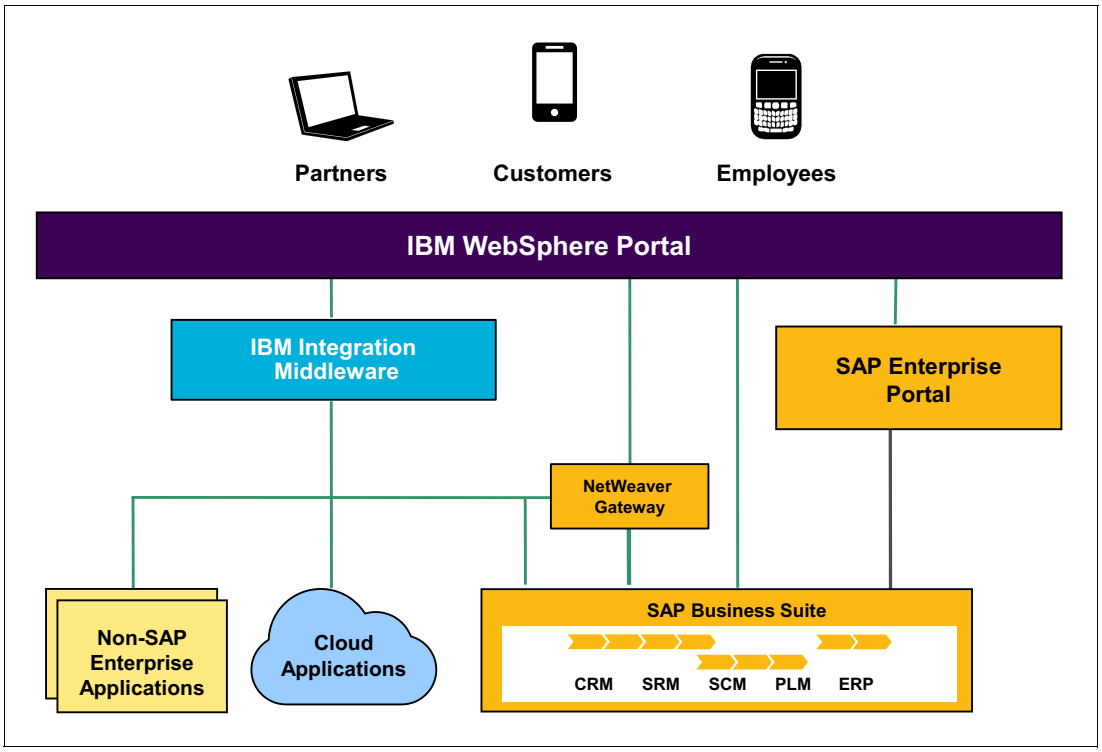

Figure 2-9 Reference portal solution for SAP

However, one approach does not meet all requirements regarding IBM WebSphere Portal and SAP integration. Different users and use cases are best served by different types of integration. The types of integration can be separated into two categories:

- Expose and reuse SAP user experience inside IBM WebSphere Portal.
- Create a new user experience to access SAP services.

The use cases for the interactions of enterprise users with SAP can be categorized into *casual* and *detailed*. The majority of employees and possibly customers will likely make casual use of SAP systems. These users need occasional access to information that originates in the SAP system. They need the information in the context of what they are doing, and do not need to know that an SAP system is involved. Casual use cases might involve a sales person looking up customer information or pricing.

Casual use cases are often best addressed by a new or simplified component that integrates with SAP at a service level. This integration option is particularly well-suited for an externally facing portal, because it can provide UI differentiation with a new user experience, one which is different from other companies that also use SAP systems.

Detailed use cases typically involve more than just simple access to SAP content. An example of a detailed use case is a sales person creating a new customer opportunity in the SAP customer relationship management (CRM) system. SAP provides a ready-to-use user experience that has been refined to meet the needs of such detailed scenarios.

This scenario is typical for intranet portals, where UI differentiation from the competition can be less important, and “on the glass” integration of pre-built SAP UI experience with a non-SAP UI can be used.

WebSphere Portal provides a pre-built integration framework that can include pre-built UI content elements from SAP Portal inside WebSphere Portal UI. This integration enables you to combine SAP and non-SAP UI content “on the glass,” and it also provides navigation integration between WebSphere Portal and SAP Portals and built-in SSO capabilities.

IBM Web Experience Factory is a model-driven rapid development tool capable of discovering SAP-provided services. It generates a rich web experience for working with SAP data based on an extensive catalog of predefined UI templates.

For more information, see Chapter 6, “Portal integration with SAP” on page 143.

2.3.7 Master data management (MDM)

Master data is key business information, such as data about customers, products, employees, and so on. Master data is usually non-transactional.

MDM solves the problem of multiple copies and sources of the same master data in the enterprise, and provides trusted and timely master data to business processes. MDM also solves the problem of multiple and inconsistent source systems for master data, by managing these business-critical information assets consistently with the highest degree of data quality.

In addition, MDM must provide the trusted master data in a timely manner, either batch-oriented or in near real-time, to all relevant business processes. SAP applications in the SAP inner ring, and the applications in the non-SAP outer ring, require access to master data.

Having an MDM solution in place to complement SAP systems adds value to SAP software by reducing the need to customize packaged SAP functionality. Without MDM, clients are often tempted to extend master data definitions within the SAP solution to suit various needs of a particular business scenario. With an MDM solution in place, the MDM platform is designed to incorporate constantly evolving master data definitions, therefore avoiding the need to customize the SAP application beyond the intended scope of the solution.

For example, the master data definition for `customer` in SAP ERP Central Component (ECC) should not contain any more attributes than those needed for the scope of the ERP package, such as financials, order fulfillment, and so on.

The master data definition for `customer` in SAP CRM should not contain more attributes than those needed to support the sales and support processes. Even within the SAP domain, SAP CRM and SAP ERP have different data models for customer and product, which further illustrates the need for an application-independent, enterprise-wide master data model for each master data entity.

MDM can also significantly reduce the burden on SAP systems by providing master data services for master data for non-SAP systems. A robust MDM solution can be much more efficient as a high-transaction, reliable system of reference for master data. As an example, consider a web commerce application that needs to display the basic customer information. Rather than having to retrieve the customer data from SAP ECC, web commerce can retrieve the information from MDM instead, therefore reducing the burden on SAP ECC.

Data quality is one of the key value propositions for adding an MDM solution to an SAP application environment. Without MDM, no assurance exists that duplicates will be properly handled, or that organizational data quality standards will be adhered to, when new master data is entered.

An estimate is that master data becomes *dirty* at the rate of 2% per month if no data quality enforcement is in place. This is particularly important for SAP applications, because completely removing master data records from SAP after they are entered is often difficult. This is especially true if operational data exists that references the master data, for example orders, invoices, and so on.

IBM InfoSphere Master Data Management delivers enterprise-scale MDM functionality that can serve both the SAP inner ring and the non-SAP outer ring, as shown in Figure 2-10. The MDM system manages master data entities, such as customer, supplier, product, and employee, providing master data with the highest degree of quality to all consumers.

In typical implementations, SAP applications hold only a copy of the master data entities (the dotted lines around the master data entities in Figure 2-10), which is managed by the MDM system. The same applies to the non-SAP applications in the non-SAP outer ring. As shown in Figure 2-10, the MDM system requires efficient integration with all of the other enterprise systems, supporting batch and real-time interfaces through the following components:

- An ESB serving both the SAP inner ring and the non-SAP outer ring
- An enterprise information integration serving both the SAP inner ring and the non-SAP outer ring

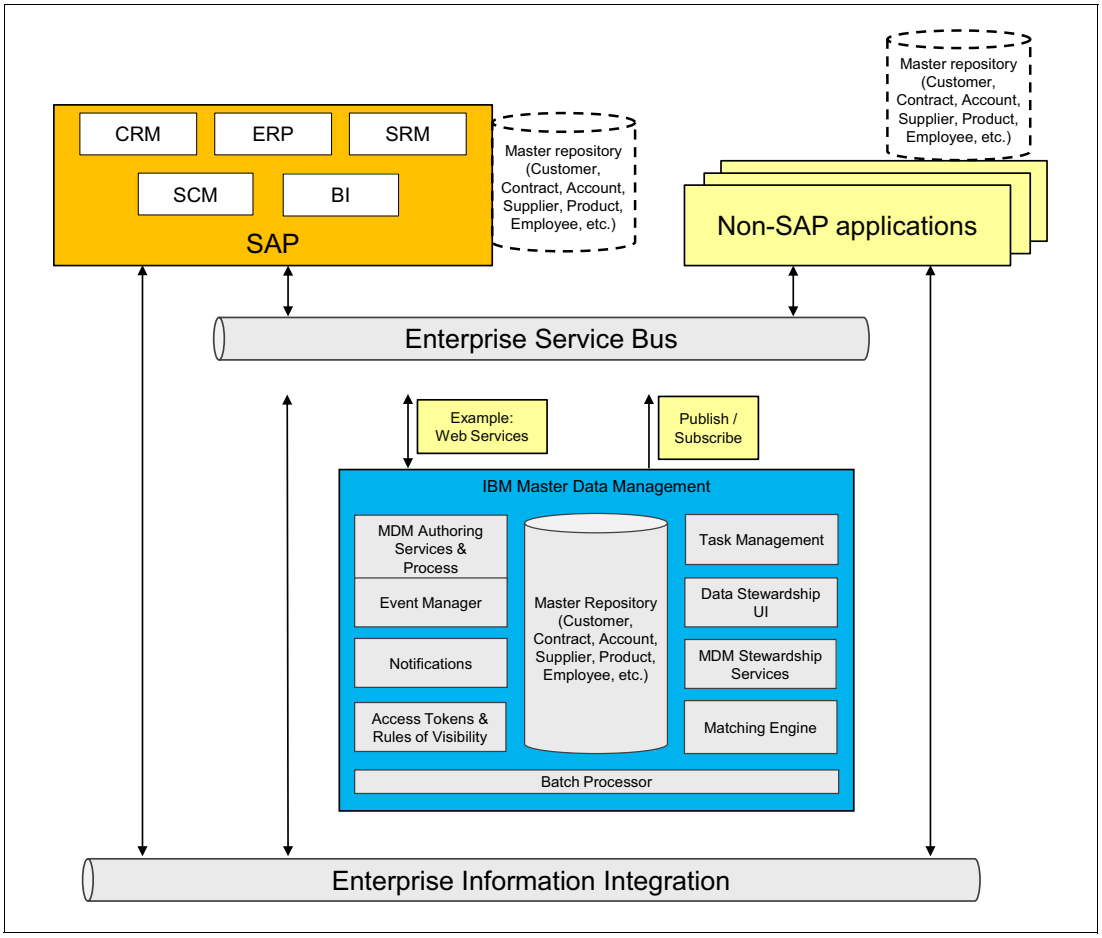

Figure 2-10 Transactional-style MDM architecture pattern for SAP in a heterogeneous enterprise

MDM is a mission-critical, enterprise-level capability for SAP and non-SAP applications, and therefore MDM is external to SAP. MDM collects and distributes master data information to consuming applications (SAP and non-SAP). The MDM platform incorporates data governance and stewardship. InfoSphere Master Data Management delivers enterprise-scale master data management functionality that can serve both the SAP inner ring and the non-SAP outer ring.

The MDM component of this reference architecture uses ESB and ETL components from already-established enterprise integration services for SAP to move data in and out of SAP systems. This covers both transactional and batch interactions.

MDM is an enterprise asset, which is typically not directly in scope for SAP adoption projects. However, MDM implementations might require a significant transformation as a result of the effect of the SAP adoption on the enterprise IT landscape. In some cases, an MDM refactoring project might need to precede large-scale SAP adoption, or run in parallel to it, but in either case it must receive adequate focus in the enterprise.

For more information, see Chapter 7, “Master data management for SAP” on page 159.

2.3.8 Enterprise content management

A large portion of the content relevant to the business in SAP systems is contained in textual or graphical form in documents, such as reports, invoices, plans, and so on. This content, which is typically referred to as *unstructured content,* is a permanent component in all business workflows. It supports and documents them in an auditable manner. All of this information is essential to support the decision processes of all SAP users.

This information must be maintained throughout its lifecycle to manage its growth and ensure legal compliance. For example, from its inception captured in paper form, converted to electronic form, throughout its way through the system with annotations, approvals, and so on, up to its disposal.

Frequently, the SAP infrastructure is not the only one that creates, operates on, and manages such information. Other systems often exist in the organization that originate documents of a similar nature, and of similar importance to the business decision-making process. Consequently, information integration between the SAP and non-SAP landscapes, to provide a unified view on all business relevant information, is vital for the operation of the enterprise.

IBM ECM software implements, and is certified for, integration through the standard SAP interfaces, such as SAP ArchiveLink, SAP Content Server, and SAP Information Lifecycle Management (ILM), and the base functionalities that SAP offers for handling this information. In addition, IBM ECM offers a spectrum of crucial extended and extendable capabilities to perform content discovery, content analytics, and case management.

The IBM Reference Architecture for SAP includes IBM ECM as part of the information architecture pattern. The ECM pattern includes the following products, each with their individual set of extended operational capabilities:

- IBM Content Collector for SAP Applications
- IBM Datacap as the capture component
- The individual ECM repositories:
  - IBM FileNet Content Manager
  - IBM Content Manager Enterprise Edition
  - FileNet Content Manager On Demand
  - IBM Tivoli® Storage Manager

IBM Content Collector for SAP Applications handles the SAP ArchiveLink and information lifecycle management (ILM) protocols to and from SAP systems, and translates them into the ECM repository-specific requests.

Non-SAP users and SAP users who choose to perform their content discovery, content analytics, and case management activities outside of the SAP GUI, can use IBM Content Navigator as the unified UI to access the federated SAP and non-SAP content, and to operate on it.

For more information, see Chapter 8, “Enterprise Content Management for SAP” on page 189.

2.3.9 Business analytics

ERP systems such as SAP provide a highly robust transactional environment that is typically optimized for transactional speed and configurability of core business functionality. ERP systems such as SAP excel at managing transactions and day-to-day business. Clients often consider the cost of deploying ERP systems to be a part of the cost of doing business.

Alternatively, business analytics solutions excel at enabling better business decisions to manage and drive business performance by providing complete visibility and fast insights into the business. IBM Business Analytics portfolio products can increase the value of SAP investments by gaining new business insights from SAP data, especially when combined with non-SAP data in a heterogeneous enterprise.

Some organizations, where users need to perform swift analysis of data from SAP solutions, require a faster-performing business intelligence solution. SAP has introduced their SAP in-memory appliance, SAP High-Performance Analytic Appliance (HANA) as a solution to help organizations analyze large volumes of detailed operational and transactional information in real-time.

IBM Business Analytics infrastructure for SAP provides near real-time in-memory analytic capabilities for SAP and non-SAP applications that do not require investment in HANA. However, if HANA is already a part of the enterprise landscape, IBM Business Analytics infrastructure for SAP provides seamless integration with HANA, as described in Chapter 9, “IBM Business Analytics infrastructure for SAP” on page 231.

The next-generation business analytic solutions from IBM help organizations of all sizes make sense of information in the context of their business. Organizations can uncover insights more quickly and more easily from all types of data, even big data, and on multiple platforms and devices.

In a heterogeneous IT landscape, SAP systems are an important source, but not the only source, of information for contextual enterprise analytics. SAP data has to be combined with business data from other enterprise systems to enable the contextual enterprise.

IBM Cognos® helps organizations to realize a greater return on their investments in SAP applications, with faster access to the data that the business needs to make better decisions. When IBM Cognos software is integrated with SAP applications, the value of SAP data is enhanced, and users gain the perspective and context needed to derive insight from SAP.

In addition, using IBM Cognos software and SAP applications together can help minimize the number of tools and duplicate content that organizations must maintain, streamline training requirements, and significantly reduce IT backlogs.

IBM Cognos TM1® is a market-leading enterprise planning software that enables organizations to collaborate on plans, budgets, and forecasts. TM1 enables users to analyze data and create models, including profitability models, to reflect a constantly evolving business environment. In addition, the integrated scorecards and strategy management capabilities help organizations monitor performance metrics, and align resources and initiatives with corporate objectives and market events.

TM1 features memory-based, multi-dimensional cube architecture. The online analytical processing (OLAP) engine driving TM1 yields excellent response times. And with multiple memory-based cubes, data is more rapidly searched, modified, and restructured than with a single-cube, disk-based structure.

Predictive analytics helps organizations to use all available data and predict with confidence what will happen next, so that you can make better decisions and improve business outcomes. IBM offers easy-to-use predictive analytics products and solutions, such as IBM SPSS®, that can use data from SAP and non-SAP sources. These solutions meet the specific needs of different users and skill levels, from beginners to experienced analysts.

The reference architecture for IBM Business Analytics infrastructure for SAP is built on the common enterprise data warehouse (EDW) that is used for both SAP and non-SAP enterprise data. This reference architecture uses SAP Business Warehouse (SAP BW) as a packaged solution.

SAP BW is an analytical, reporting, and data warehousing solution produced by SAP. SAP BW is a packaged solution that includes SAP delivered extractors, data models, and queries (also known as *Business Content*). In this reference architecture, the scope of SAP BW is limited to operational reporting on SAP Business Suite data. All deeper analytics and all analytics at an enterprise level should be based on the EDW.

Figure 2-11 on page 34 provides an overview of IBM Business Analytics integration capabilities for SAP.

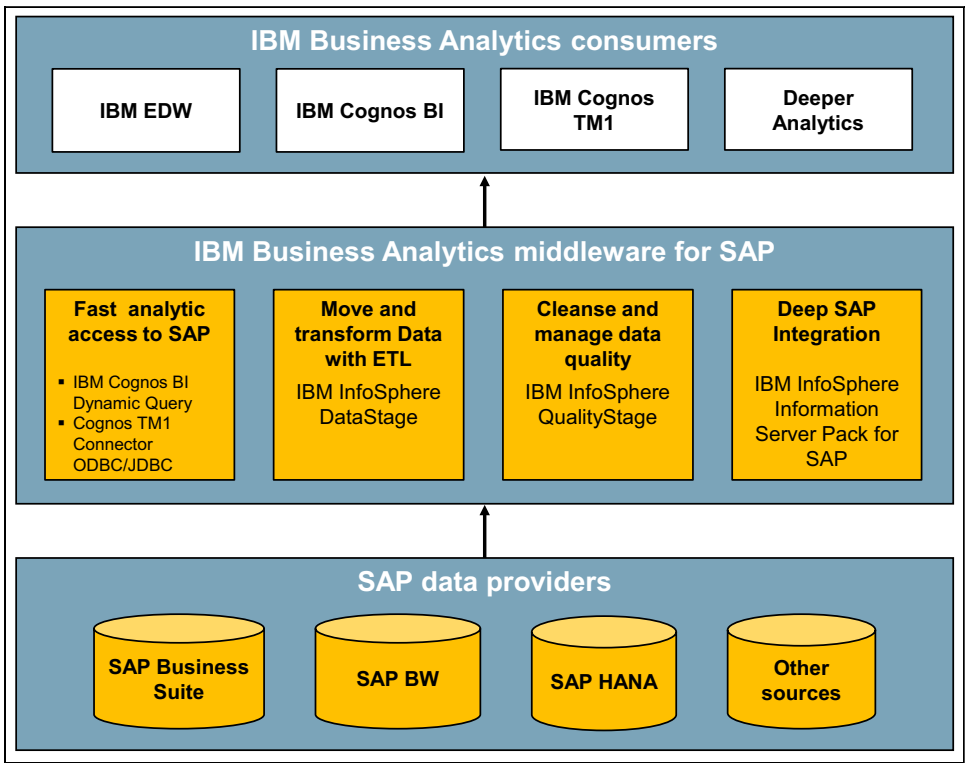

Figure 2-11 Data warehousing and business analytics for SAP in a heterogeneous enterprise

Figure 2-11 shows that IBM middleware has capabilities to perform these tasks:

- Connect to any SAP source.
- Deeply integrate with SOA data.
- Provide near real-time SAP data replication into analytics solutions.
- Provide more structured data extraction and cleansing capabilities.

The IBM InfoSphere Information Server is a key component that encapsulates best-in-class integration tools to collect metadata, and to manipulate or assess data before integration with consumer BA applications. SAP integration is based on using SAP-certified integration interfaces:

- Open Database Connectivity (ODBC)
- Database Connectivity (JDBC)
- OLAP Business Application Programming Interface (BAPI)
- SAP Remote Function Call (RFC) through Advanced Business Application Programming (ABAP) business functions

For more information, see Chapter 9, “IBM Business Analytics infrastructure for SAP” on page 231.

2.3.10 DevOps for SAP

To realize the full value of an investment in SAP solutions, teams need to enable continuous delivery of change across their SAP landscape and the dependent non-SAP technology. More than 80% of SAP implementations are heterogeneous in nature, and delivery teams require a single integrated solution across both the SAP and non-SAP components to achieve an end-to-end centric transformation program.

Traditional approaches to SAP development, integration, management, and delivery rely largely on manual processes. This approach is error-prone and time-consuming, delaying the response to business needs and IT change requests. Silos of processes, projects, people, and tools can make it difficult to gain a clear view of systems and processes, and to align business and IT requirements. For organizations looking to improve the value delivered by SAP solutions, a more unified and flexible approach is needed.

IBM believes that the key to addressing these challenges lies in the IBM DevOps solution, a contraction of *development* and *operations*, the two teams that deliver the core IT services in an organization. By bringing these teams closer together, and applying agile principles across the entire SAP delivery lifecycle, the IBM DevOps solution promotes continuous innovation, feedback, and improvement. This approach enables you to accomplish the following goals:

- Help development *and* operations teams boost productivity.
- Improve the value delivered to users.
- Accelerate delivery.
- Reduce the cost and effort of IT change.

IBM has the tools and technology to achieve continuous delivery, and to reduce the cost and risk of managing changes to the SAP landscape. IBM calls this set of tools and technologies *IBM DevOps for SAP*.

Figure 2-12 shows how the IBM DevOps solution complements and extends SAP, using a set of pre-built and pre-integrated components to provide an effective lifecycle management solution in a heterogeneous enterprise.

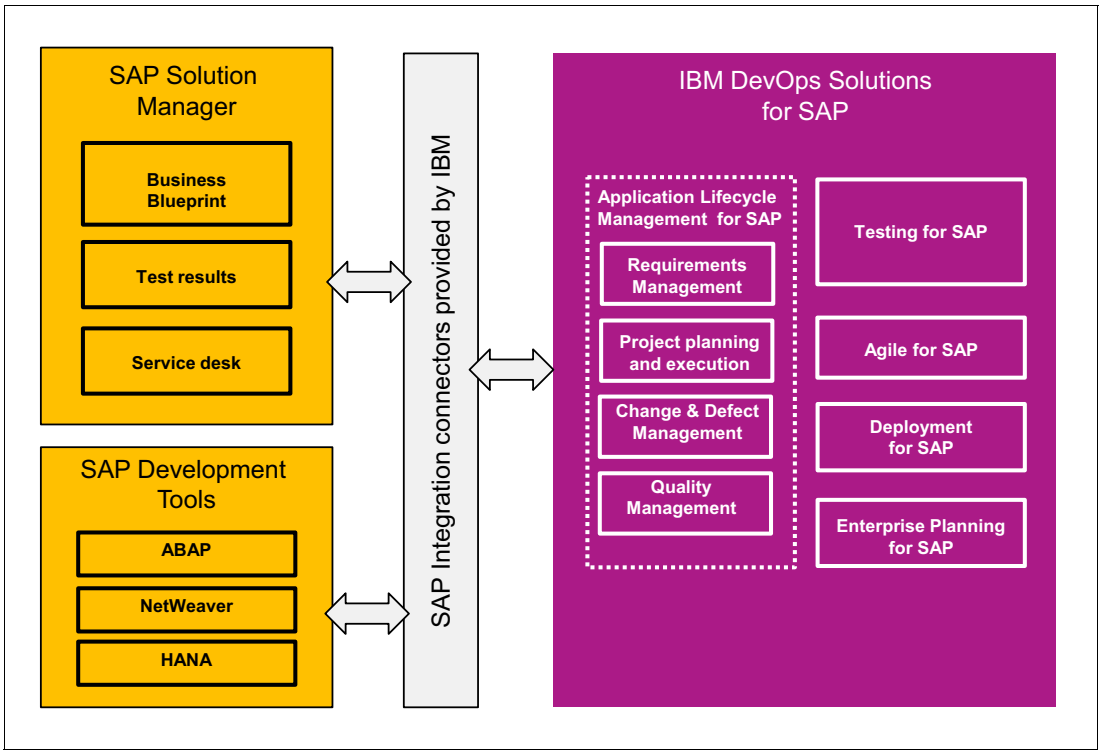

Figure 2-12 the IBM DevOps solution extends SAP application lifecycle tools in a heterogeneous enterprise

Requirements management

SAP Business Blueprint is a form of business requirements definition that is typically used in SAP projects. IBM requirements management solutions enable you to effectively manage any

form of business requirements definitions. IBM requirements solutions fully support the SAP-mandated business requirements management process (blueprinting), based on SAP Solution Manager, that is used to manage all SAP-related Blueprint items using a traditional SAP approach.

IBM requirements management extends SAP, and should be used to manage all the requirements related to non-SAP components of an overall SAP-centric solution. Alternatively, SAP Solution Manager can be used to document the structure of SAP Blueprint (business process hierarchy).

However, all of the SAP Blueprint content is managed in IBM requirements management as structured data, which provides further structure and decomposition to SAP-controlled Blueprint business process hierarchy. For example, it provides more detailed requirements decomposition for additional technical components that must be included in SAP projects, such as (WRICEFs).

This approach provides more structured, traceable, and integrated management of the actual contents of SAP Blueprint when compared to the traditional all-SAP approach, whereby such requirements are managed manually as basic document attachments.

Overall, the guideline from this reference architecture is to use SAP Solution Manager to capture business process hierarchy for SAP Blueprint, mirror it in IBM requirements management, and use the latter to manage actual requirements contents as structured data.

Project planning and execution

In heterogeneous enterprise environments, SAP delivery projects often involve work affecting multiple disparate applications and systems. With the IBM solution for project planning and execution, SAP project teams can use a single tool for project and iterations plans, work item tracking, deliverable activity management, and release planning.

IBM Rational Team Concert™ covers multiple areas of functionality common in all variations of project planning and execution. IBM Rational Team Concert can be used to manage SAP and non-SAP projects in a unified way. This approach enables teams to plan and run projects based on end-to-end business processes, and to coordinate all changes and release deliveries across the different applications and systems.

IBM Rational Team Concert is also used as the collaborative project hub to track all work, control project governance, and identify gaps in work items.

Change and defect management

Change management is a key component in the solution lifecycle governance. Effective change request management requires a common information repository that all team members can access.

In this reference architecture, IBM Rational Team Concert is used to define and govern changes throughout the project lifecycle. IBM Rational Team Concert provides a single change management platform for both SAP and non-SAP change requests.

Change requests that are managed by IBM Rational Team Concert are not limited to source code changes; they can manage other changes in the SAP solution, for example UI changes. Complex changes can be delivered to all affected SAP and non-SAP applications and middleware systems in a synchronized fashion.

Defect management is typically considered a variation of change request management. IBM Rational Team Concert supports an integration with the SAP Service Desk so defects and enhancements that need to be coordinated with the SAP help desk can be automatically

managed and tracked. For example, a specific defect can be forwarded to the SAP Service desk. The defect submission form is populated with live data from SAP Service Desk.

Quality management

Quality management consists of three main functional areas:

- Test planning
- Test execution
- Test reporting

In addition, defect management is typically viewed as an extension to quality management, although within the context of application lifecycle management it is usually captured as a specific variation of change request management.

IBM Rational Quality Manager is one of the testing and quality solutions endorsed by SAP. It provides extended capabilities beyond the SAP solution built into SAP Solution Manager. IBM Rational Quality Manager is used for SAP and non-SAP-centric quality management, and specifically for test planning and reporting.

IBM provides pre-built integration connectors with SAP that enable IBM Rational Quality Manager to be integrated with SAP Solution manager. With this approach, you can automatically map elements in the SAP business Blueprint to test plans and test cases. The test results from IBM testing tools are automatically synchronized back into SAP Solution Manager at the appropriate level within the Blueprint. The Blueprint becomes a general business-focused container for the overall test architecture.

Test execution

IBM and IBM Business Partners provide an extensive set of testing tools for SAP functional, integration, performance, and security testing. For more information, see Chapter 10, “DevOps for SAP” on page 249.

Agile for SAP

Agile software development and delivery is mainstream for custom application development. It is increasingly being used to successfully develop and deliver packaged applications, including SAP applications and changes.

The IBM Rational Agile for SAP offering (part of the IBM DevOps solution) provides ready-to-use, customizable agile and lean planning, change, work, and delivery management and execution method support for SAP applications. The core of this solution is IBM Rational Team Concert, a market-leading agile project planning, change, defect, and delivery management solution.

IBM Rational Team Concert provides pre-built integration capabilities with the Eclipse-based SAP integrated development environments (IDE) such as SAP NetWeaver IDE and SAP HANA IDE.

The following list describes the key benefits of IBM Rational Agile for SAP offering for SAP projects:

- Improve SAP developer and team productivity by enhancing traditional SAP development with proven advantages of agile methods.
- Enable better decisions based on real-time, transparent visibility into SAP Agile delivery and maintenance projects.
- Accelerate agile adoption and results using pre-configured and customizable agile (or others) method definition and automated enactment.

Continuous release and deployment for SAP

When businesses run on SAP, their success depends on how well they manage IT change. An inability to continuously deploy can negate the investment that organizations make in developing and implementing change into business processes and integrations.

Traditional approaches to SAP deployment rely largely on manual processes and spreadsheet management. These approaches can be error-prone and time-consuming, delaying the response to business needs and IT change requests. SAP application landscapes are most often composed of many different applications that are interdependent and can often involve both SAP and non-SAP applications and components.

The IBM Rational Continuous Deployment and Release for SAP solution can be used to coordinate the deployment and release of SAP and non-SAP landscapes. Specifically, it provides the following capabilities:

- Plan, coordinate, orchestrate, and manage releases of integrated application deployments.
- Automate the deployment of required non-SAP configured middleware.

These capabilities lead to quicker releases, improved communications, and lower risk associated with the software releases.

Enterprise planning for SAP

Successful IT transformation investments in SAP establish an enterprise architecture to provide organizations with a big picture, and the framework to select the best overall route for SAP adoption in a heterogeneous enterprise.

By effectively incorporating SAP software into the enterprise architecture, organizations gain better insight into their overall technology plans. With an enterprise architecture in place, organizations can ensure that the SAP business model as defined is the business model that they implement, and can guide SAP investments to run outcomes that support their corporate strategy.

IBM Rational System Architect is an enterprise architecture tool that enables organizations to effectively analyze and plan their business and technology architecture. IBM Rational System Architect and IntelliCorp LiveCapture for SAP Solution Manager helps organizations understand how their SAP systems map to their overall architecture, what happens when architectural changes are made, and how the SAP environment can be used across the enterprise.

The following list describes key capabilities of Rational System Architect:

- Automated synchronization of models from SAP Solution Manager into Rational System Architect
- Visualization and integrated view of SAP projects, blueprints, and landscapes in the context of the enterprise architecture, and of business processes, data, organization, and roles in a business-process context
- Comparisons of *as-is* and proposed *to-be* solutions

For more information, see Chapter 10, “DevOps for SAP” on page 249.

3

Enterprise integration services for SAP

A key observation from Chapter 2, “IBM Reference Architecture for SAP” on page 9 is the fact that any major enterprise has a heterogeneous application landscape composed of SAP and non-SAP applications. The purpose of this chapter is to introduce enterprise integration services and explain how the various applications are efficiently integrated. Enterprise integration services for SAP include both near real-time transactional and batch-oriented integration techniques.

For most SAP solutions, integration activities can be divided into two major categories:

- Initial data load (also referred to as *conversion* in many SAP projects)
- Ongoing integration (also referred to as *interfaces* in SAP projects)

As the name suggests, initial data load involves the one-time activities required to move large volumes of data from existing systems into the new system (in this case, SAP and related application components) before going live. The ongoing integration involves all of the interfaces between existing systems and the new system performed on a regular basis after the new system is live.

The ongoing integration encompasses a wide range of integration requirements and approaches: Batch versus transactional, synchronous versus asynchronous, service-oriented versus message-oriented, and so on.

This chapter provides a coherent architectural blueprint for both ongoing integration with SAP applications and initial data load into SAP applications.

This chapter includes the following topics:

- 3.1, “Introduction to enterprise integration services for SAP applications” on page 40
- 3.2, “Architecture goals” on page 41
- 3.3, “Scenarios and patterns for ongoing integration with SAP” on page 42
- 3.4, “Architecture overview of ongoing integration with SAP” on page 51
- 3.5, “Architecture components of ongoing integration with SAP” on page 52
- 3.6, “Initial data load” on page 66

© Copyright IBM Corp. 2014, 2015. All rights reserved.

3.1 Introduction to enterprise integration services for SAP applications

Integration takes place between interfaces. Therefore, to integrate with SAP applications, you must first understand which interfaces SAP applications offer for integration.

Figure 3-1 shows a conceptual architecture overview for operational SAP applications, such as SAP Enterprise Resource Planning (ERP) or SAP Customer Relationship Management (CRM). The figure shows that the Web Application Server run time is the common runtime infrastructure for the SAP applications. The Web Application Server (Web AS) is implemented in the SAP proprietary programming language, Advanced Business Application Programming (ABAP).

Older versions of the Web Application Server, for example versions 6.20 and 6.40, were the run time for the SAP R/3 systems. Newer versions, starting with version 7.00, became the core of the SAP NetWeaver platform.

As shown in Figure 3-1, the key interfaces to integrate with SAP applications are Business Application Programming Interface (BAPI), Intermediate Document (IDoc), ABAP (for extract) and SAP enterprise services. All of these interface types are explained in “Endpoint connectivity” on page 44.

All of these interfaces operate on logical models at the application level, because a key design decision by SAP has been to decouple the applications from the underlying databases, such as IBM DB2, Oracle, and so on. Therefore, integrating with SAP applications on common database interfaces, such as Java Database Connectivity (JDBC), is not possible.

The only exception to this rule is the recent addition of SAP High-Performance Analytic Appliance (HANA) as persistency where access through Open Database Connectivity (ODBC) is enabled.

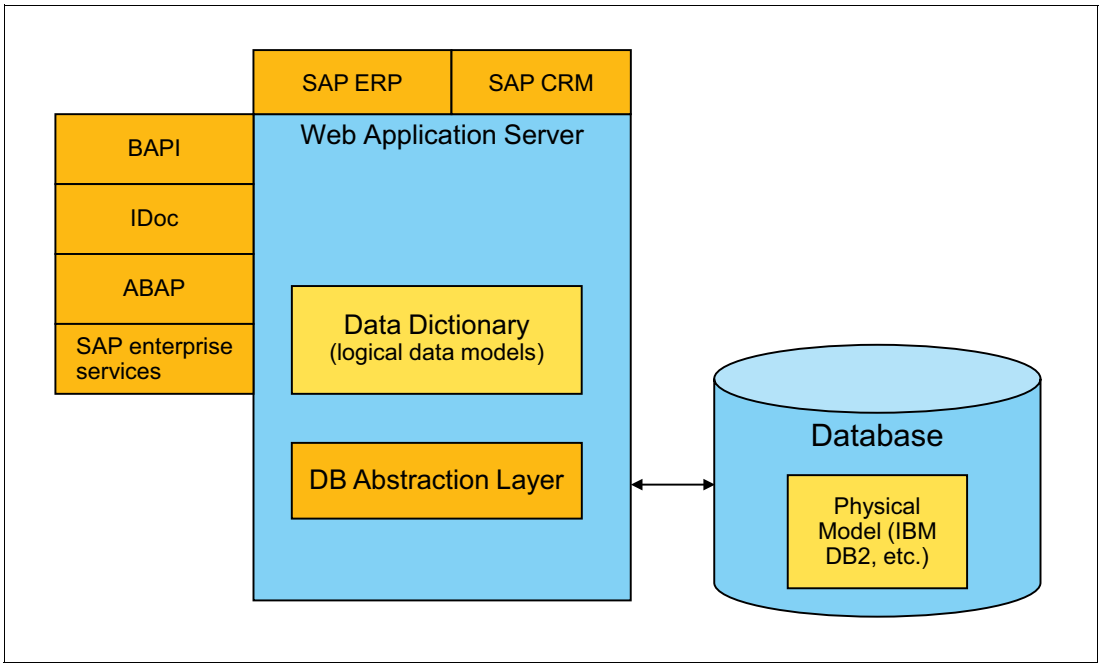

Figure 3-1 SAP integration interfaces overview

Figure 3-1 on page 40 does not show additional interfaces for analytical SAP applications such as SAP Business Information Warehouse (BW) and SAP Business Intelligence (BI), which support interfaces such as Open Hub.

For many years, IBM has provided SAP certified connectivity using all of the integration mechanisms with SAP applications described in this section, enabling both transactional and batch integrations.

3.2 Architecture goals

This section introduces the architecture goals for enterprise integration services for SAP applications.

3.2.1 Align enterprise integration services with SAP implementation methodology

In the SAP implementation methodology, two major integration categories are captured as two separate types of work products:

- Conversion Definition Document (CDD)
- Interface Definition Document (IDD)

The details of these definition documents are outside the scope of this document. The focus of this document is the integration between SAP and non-SAP systems.

The IDDs identified during the blueprint phase of the SAP methodology do not make a distinction between the different types of interactions, such as transactional versus batch.

IBM software provides technologies to address the different integration requirements. The goal of the architecture blueprint described in this chapter is to provide a coherent coverage for both transactional and batch integration aspects, and to provide guidelines for selecting appropriate IBM integration middleware based on SAP integration requirements.

3.2.2 Use best-in-class technologies for custom integration development

SAP has several available technologies that provide integration capabilities. In some cases, the technologies include pre-built integration modules that can be configured for a particular environment, and used without the need for any custom development.

However, customizations made within the SAP applications limit the extent to which the pre-built integration capabilities can be used without modification. Additionally, many of the SAP integration technologies require custom development to provide the integration logic and do not provide any special capabilities with regard to SAP application components. Even within the SAP inner ring (described in Chapter 2, “IBM Reference Architecture for SAP” on page 9) the effort and approach to develop new integration logic with IBM software technologies is the same as when using SAP integration technologies.

In the cases where SAP integration components can meet 90% or more of the integration requirements with pre-built functionality, the SAP components should be used in the SAP inner ring. However, if more than 10% of the required functionality must be provided through custom development, the integration should be implemented using best-in-class, IBM technologies, even if the integration is performed within the SAP inner ring.

3.2.3 Minimize costs of integration for non-strategic systems

For large SAP application deployments, a transition period usually exists during which the new SAP system begins to assume functionality previously provided by existing applications. During the transition, the new functionality is not yet fully rolled out. As a result, two types of integration requirements exist for transition architecture:

- Non-SAP applications that are non-strategic in the long term, but must be integrated with SAP applications and remain operational until the new system can fully assume the business functions
- Non-SAP applications providing strategic functionality that will not be incorporated into the SAP solution, and that need to be permanently integrated with the new system

Strategic non-SAP applications should be integrated with SAP applications according to established standards and leading practices. Integration costs can be justified with longer-term benefits in maintenance cost reduction, flexibility, adherence to strategic objectives for the enterprise, and so on.

However, for those non-SAP applications that are expected to be decommissioned after the new system is in place, there is often little justification for making further investments. Depending upon the status of the application, making additional changes might not even be possible.

As a result, the approach to integrate non-strategic existing applications can require changes in either the application or the integration layer. In this case, the least expensive integration approach overall should drive the changes. The integration of non-strategic applications is usually short-lived after the transition is completed. Therefore, limiting the expense and effort put into it is important.

3.2.4 Loosely coupled applications

To the greatest extent possible, the systems should be integrated so as to minimize hard dependencies between the application components. Loose coupling enables one component to be replaced with little to no effect on the other components. It enables components to be reused by other consumers. It also isolates systems to improve resiliency and enable for components to be defined and developed independently.

3.2.5 Use open, well-established standards

Standards-based implementation of integration logic greatly reduces the chance of vendor lock-in. It also increases the availability of tools that easily connect to the solution and provide support for development, monitoring, and testing. As a practical point, when attempting to move two proprietary systems to a common format or approach, selecting an open, industry standard can often provide a neutral option that prevents the *mine versus yours* arguments.

3.3 Scenarios and patterns for ongoing integration with SAP

To create an efficient integration layer that can be operated and maintained in a cost-effective manner, identifying a standard set of integration patterns that can be applied for common integration scenarios is important. The following sections provide guidance on how to classify the interfaces based on the integration requirements. The appropriate integration scenario and preferred integration patterns for those scenarios are identified.

3.3.1 Identifying integration scenarios

The information in this section can help you identify the appropriate integration pattern.

Interface characteristics

The first step in selecting an integration pattern is to identify the characteristics and requirements of the integration itself. Many of the characteristics can be applied independently of each of the systems involved in the overall interaction, which causes variations within the patterns. For example, the originator uses a one-way message exchange pattern, although the destination uses a request-response pattern.

Interaction style: Transactional or batch

The first major division among the various interfaces is the interaction style: Does the interface represent a single business transaction and business object, or does it represent a collection of multiple business objects? The answer to this question results in either a transactional or a batch integration style.

Remember: The term *transactional* does not refer to requirements dealing with the transactional integrity of the interface. Rather, it indicates that the interface represents a single *business transaction*. This identification should be made solely on the contents of the message, and not on other outside factors, such as the frequency of the interface or the planned transport.

The following examples help to identify the interaction style:

- A query object containing a set of attributes is submitted with a response that contains multiple results: `Transactional`.
- A collection of existing customer numbers are sent to a customer translation service to return a collection of proper customer identifiers: `Batch`.
- A message flow sends a business event containing the current message processing context to the event infrastructure in a fire-and-forget operation: `Transactional`.
- Every five minutes, a business application is triggered to collect all of the changed records and send them to the destination system: `Batch`.

Message exchange pattern: Request-response or one-way

The message exchange pattern identifies the interface as either `request-response` or `one-way`. If the sender of the message expects an answer back from the request that was sent, regardless of the timing of the answer, it is a `request-response` exchange. If a message is sent without a business response, it is `one-way`.

Message response timing: Synchronous or asynchronous

The response timing indicates whether the interaction is `synchronous` or `asynchronous`. Generally, asynchronous interactions are preferred over synchronous interactions if an immediate response is not required, because it gives greater flexibility for managing the load.

Interaction effect: Read-only or change

Two options exist for the effect of the interaction: `Read-only` or `change`. `Change` applies to interfaces that make changes to the back-end system. `Read-only` applies to interfaces that are simply running a query to read information.

Message handling: Independent, grouping, or sequential

Message handling refers to how to handle the messages as they arrive in the middleware software. The messages are either independent of one another, part of a larger group, or require sequential handling:

Message handling: Independent
The most common and simplest situation for the message sources is independent processing: The messages are processed independently of one another as they arrive.

Message handling: Grouping
Grouping means that the messages are part of a group, and processing cannot begin until all of the messages arrive and they can be processed together.

Message handling: Sequential
Sequential processing indicates that the order of the message processing is dependent either on the order in which the messages arrive or on some sequencing information found within the messages.

Destination cardinality: Single or multiple

Destination cardinality refers to the number of destination systems (service providers) that are involved in the interaction: A *single* system or *multiple* systems.

If the cardinality is *multiple*, several subcategories exist:

Multiple destinations: Routing
Although there can be multiple possible destinations, any particular message is sent to one and only one of the destinations, based on some set of business rules. The business rules might evaluate the business elements within the message in addition to the timing and context of the operation.

Multiple destinations: Aggregation
In this case, all of the destinations are involved in the interface. Each of the service providers is started independently, and the results are compiled together into a single response message.

Multiple destinations: Broadcast
For broadcast, often called *publish/subscribe*, the received message is sent to any number of the potential destinations, or none at all. Each of the destination systems has independently established business rules defining which messages they are interested in receiving. When the message is received, each of the rules is evaluated and the message is sent only to the matching destinations.

Multiple destinations: Sequential dependency
For sequential dependency, multiple destinations are dependent on one another in some way. For example, a message is sent to System A and after a successful response has been received a message is sent to System B, then System C, and so on. In some cases the response from one system is used to create the request for the next system.

Endpoint connectivity

Several options are available for connecting application endpoints with the enterprise integration service middleware components. Each of the connectivity options varies in terms of availability and capabilities, and some options are preferred over others.

SAP connectivity

Table 3-1 summarizes SAP endpoint connectivity options that are available for different interface styles.

Table 3-1 SAP connectivity options

Interaction type	Transactional (individual)	Batch
Synchronous	► SAP Enterprise Services ► Remote Function Call (RFC)	N/A
Asynchronous	► SAP Enterprise Services ► IDoc	► SAP Enterprise Services ► IBM ABAP Stage ► IDoc ► File system

The following list briefly describes the SAP endpoint connectivity options and their main characteristics:

- SAP Enterprise Services. SAP Enterprise Services are a set of web services definitions provided out-of-the-box by SAP with the following attributes:
  - Based on Web Services Description Language (WSDL) and SOAP web services standards
  - Based on SAP global data types
  - Modeled in SAP Enterprise Service Repository (ESR) using business objects, process components, and the SAP enterprise model
  - Published in the SAP Service Registry (SR)
  - Ensured availability and functional correctness
  - Support for Web Service Reliable Messaging (WSRM) for select services

  SAP Enterprise Services support synchronous and asynchronous transmission styles, and can be used for both transactional or batch interaction styles. SAP Enterprise Services should be evaluated for use whenever one is available for a particular business function.

  The evaluation should be made based on how well the Enterprise Service matches the defined interface characteristics, and how well the Enterprise Service addresses the particular non-functional requirements. However, the catalog of available services is rather limited, so the opportunity for using them might be small.

  Tip: If an SAP Enterprise Service for a particular business function is not available, it is *not* considered a good practice to create a custom web service directly within the SAP environment. Instead, the services should be developed using best-in-class tools for developing and managing services. This approach enables you to expose the SAP function through the enterprise service bus (ESB) in a standard manner across the enterprise.

- Remote Function Calls. RFCs provide real-time interactions with SAP systems in a request-response message exchange pattern. The RFCs include ready-to-use function modules called BAPIs, in addition to custom function modules. RFC operations block the thread within the SAP systems while the operation is performed against other function modules and the database. As a result, use RFCs to perform transactional (individual business operations) and synchronous interactions with SAP systems.
- Business Application Programming Interface. BAPIs are well-defined external SAP interfaces providing access to processes and data in SAP business application systems. BAPIs are designed to be started by systems external to SAP using the RFC mechanism.
- Advanced Business Application Programming. ABAP is a proprietary programming language from SAP. This interface is used by IBM integration tools to discover data objects in SAP systems, such as SAP tables. It automatically generates and deploys into the SAP system RFC-enabled ABAP code modules that are subsequently used for extracting data from SAP tables. This mechanism is used only to *extract* SAP data and never to update SAP data, because updates to SAP data must be done only through SAP business logic.

  IBM ABAP Stage is a component of IBM InfoSphere Information Server Pack for SAP Applications (SAP BW) where ABAP logic is generated by the tool to perform the data extraction logic from SAP (typically from SAP tables). IBM ABAP Stage should be used in batch interfaces where data must be extracted from SAP and sent to non-SAP systems, and where ready-to-use operations are not already available in SAP to perform the data extraction, that is, situations where custom development would otherwise be required.
- Intermediate Document. Within SAP systems, IDocs provide a particular hierarchical message format that can be posted to SAP asynchronously using standard RFC transports, including mechanisms for transaction management and queuing.

  IDocs typically represent data of an SAP business object or array of business objects, for example, a purchase order. An IDoc should be used in any one-way interface with SAP systems (batch or otherwise). Additionally, two IDocs (one inbound and one outbound) can be used to implement asynchronous request-response message exchange patterns.
- File system. In some rare cases, particularly when using custom existing SAP function modules, it is necessary to provide information to the SAP system from the file system. A typical setup includes a local file system or Network File System (NFS) mount on the SAP application side. Some middleware component (ESB, extract, transform, and load (ETL), or Reliable File Transfer (RFT) at a minimum) is between the file system and the non-SAP system, so that the non-SAP systems do not put files directly on the SAP file system.

Non-SAP connectivity

Table 3-2 is a summary of non-SAP endpoint connectivity options available for the various interface styles.

Table 3-2 Non-SAP connectivity options

Interaction type	Transactional (Individual)	Batch
Synchronous	► Web services, such as Hypertext Transfer Protocol (HTTP) and HTTP Secure (HTTPS) ► Representational State Transfer (REST) or RESTful, services	N/A
Asynchronous	► Web services, such as IBM MQ and Java Message Service (JMS) ► Web services (HTTP / HTTPS) ► Message-based IBM MQ	► Web services (IBM MQ / JMS) ► Message-based IBM MQ ► RFT ► Secure File Transfer Protocol (SFTP) ► Database

The following list describes the key characteristics of each endpoint connectivity option:

- Web services (HTTP, HTTPS). HTTP-based web services are perhaps the simplest way to perform integration. The interface specification is defined using WSDL, and SOAP defines the structure of information exchanged.

  Much of the work that is required to perform the integration using web services is provided by available tools. By adopting additional web service standards, support can be added for other capabilities, such as security (WS-Security), transactional integration (WS-Transaction), and reliable messaging (WS-RM).

  HTTP-based web services are traditionally used in synchronous request-response or one-way interactions. They can also be used in an asynchronous fashion using either `Request with Callback` or `Request with Polling` patterns.

  In the `Request with Callback` pattern, the response is sent by making a separate call from the service provider (or ESB) back to the service consumer. In the `Request with Polling` approach, the response from the initial request is a ticket that identifies the request. The consumer makes subsequent calls using the provided ticket to retrieve the response if it is available.

- RESTful Services. RESTful services use HTTP methods and Uniform Resource Identifiers (URIs) to perform create, read, update, and delete (CRUD) operations. RESTful services are commonly used with mobile and Web 2.0 applications.

- Web services (IBM MQ and JMS). IBM MQ (IBM WebSphere MQ) based web services have WSDL definitions for the interface and use SOAP over IBM MQ transports. These services are ideal for interfaces that operate in an asynchronous fashion, both batch and individual interaction styles.

- Message-based IBM MQ. Message-based IBM MQ (as opposed to service-based IBM MQ) is the traditional approach to using IBM MQ where the messages placed on the queue might be in any format and where the general context for the message is provided by the queue upon which the message arrives.

- Reliable File Transfer. RFT is the transport protocol for moving files around the infrastructure in a centrally controlled manner. It is a replacement for traditional approaches such as FTP and NFS. RFT can be deployed transparently to the source and destination systems. The systems deal with files on the local file system, and RFT does the work required to get the files where they need to go.

- File Transfer Protocol. FTP (and related technologies such as SFTP and FTPS) enables systems to send and receive files in a point-to-point manner. The connectivity information and the security certificates must be managed by each endpoint application.

- Database. Appropriate only in special situations, direct database interactions can be used in batch interfaces to retrieve or store large volumes of information. The challenge with using direct database interactions is that it almost always requires detailed application logic within the middleware, particularly if the connection is made to the application's operational database as opposed to an information warehouse or staging tables.

3.3.2 Common integration patterns

A large number of possible combinations of integration interface and endpoint connectivity characteristics constitute various integration patterns. This section describes the most commonly used integration patterns.

Simple batch

The simple batch integration pattern shown in Figure 3-2 uses the batch capabilities of the ESB software, for example, IBM Integration Bus, to satisfy simple batch requirements.

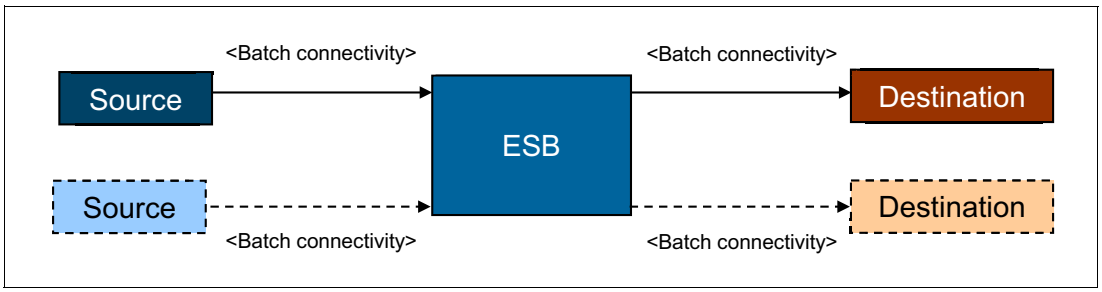

Figure 3-2 Integration pattern: Simple batch

A simple batch integration scenario includes the following characteristics:

- Independent message handling, which indicates that no grouping or sequencing requirements exist.
- Basic mediation requirements for the middleware, for example, transformation and basic field translation.
- Required connectivity does not include database interactions, nor does it require complex extract logic from SAP using IBM ABAP Stage.
- Modestly sized batch messages, for example, less than 25 megabytes (MB).

Simple batch integration can have a single source or multiple sources, depending upon the nature of the data and the interface (the dashed lines in Figure 3-2 indicate optional components). For example, either a single system provides a particular set of business data, or multiple systems do. Similarly, simple batch integration can have a single destination or multiple destinations. However, for an interface to be considered simple batch, there should be no dependency between the sources and destinations.

Complex batch

The complex batch integration pattern shown in Figure 3-3 applies to more complicated batch integration scenarios that require a full-featured ETL batch integration platform, such as the InfoSphere Information Server component IBM InfoSphere DataStage®.

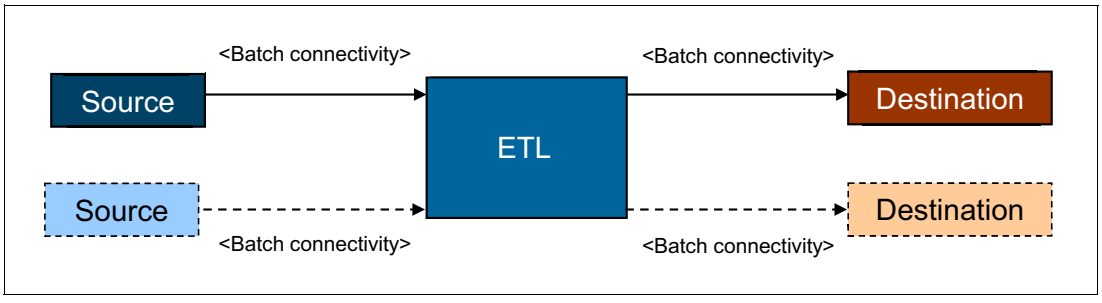

Figure 3-3 Integration pattern: Complex batch

A complex batch integration scenario includes one or more of the following characteristics:

- Grouping or sequential message handling.
- Complex mediation requirements, for example, cleansing, staging, complex transformation, complex field translation.
- Connectivity with source or destination systems uses database interactions and complex extract logic required from SAP using IBM ABAP Stage.
- Large messages sizes (greater than 25 MB).

As with simple batch integration, the cardinality of the sources and destinations is one-to-many for complex batch integration, meaning that either side can have multiple systems. Unlike simple batch, complex batch can include dependencies between the various systems or the messages of a single system.

One-way transactional

In this integration pattern, as shown in Figure 3-4, one or more service consumers send a one-way message that represents a single business transaction to the service provider.

Optionally, multiple service providers can be accommodated, assuming the logic follows either the `routing` or `broadcast` approach for multiple destinations.

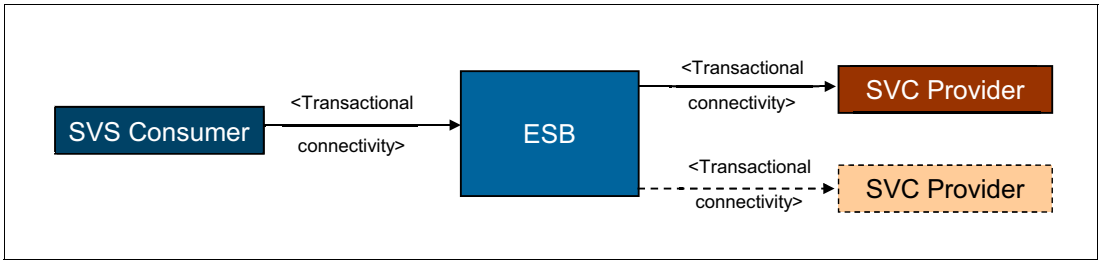

Figure 3-4 Integration pattern: One-way transactional

Synchronous transactional request-response

In this integration pattern, as shown in Figure 3-5, messages are sent synchronously from the service consumer through the ESB to the service provider.

Optionally, multiple service providers can be supported if the ESB selects one of the service providers and routes the message to the appropriate endpoint.

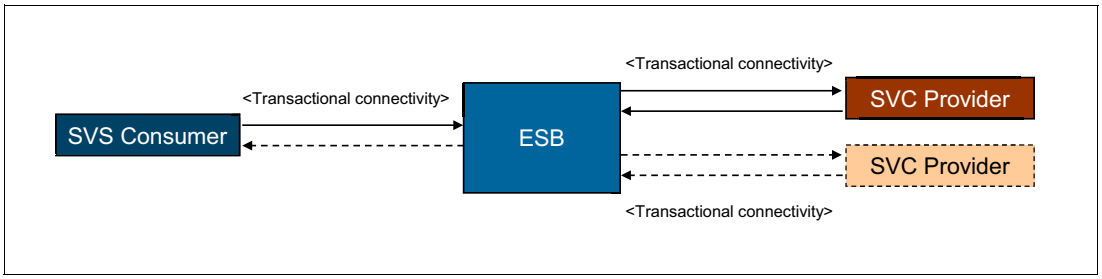

Figure 3-5 Integration pattern: Synchronous transactional request-response

Asynchronous transactional request-response

If the interaction can be performed asynchronously, the asynchronous transactional request-response pattern shown in Figure 3-6 should be used. The asynchronous interaction on the consumer side can support several implementation approaches simultaneously, for example IBM MQ web services, HTTP web services with callback or polling, and so on. Optionally, multiple service providers can be supported if the ESB selects one of the service providers and routes the message to the appropriate endpoint.

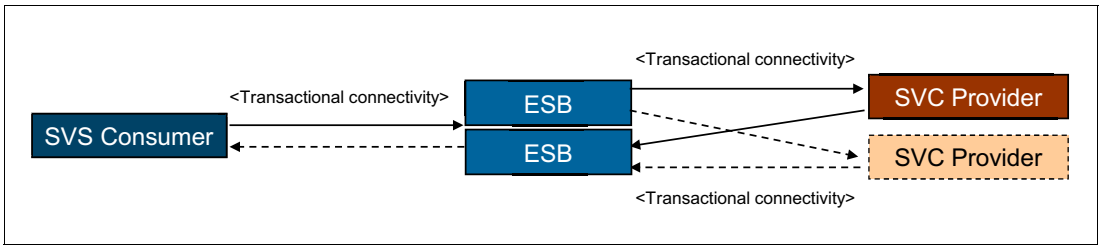

Figure 3-6 Integration pattern: Asynchronous transactional request-response

Figure 3-6 represents the ESB as two different boxes to indicate that the interaction is performed in a separate thread with a different connection. In this case, the response message is assumed to be performed in a different context from the request. This approach also assumes that the context is available in the response message or can be looked up within the ESB, for example, reply-to address of the service consumer.

Simple orchestration

Figure 3-7 shows the simple orchestration pattern. This pattern describes an interface that involves coordination and management of multiple destination services, particularly the aggregation or sequential dependency interface characteristics.

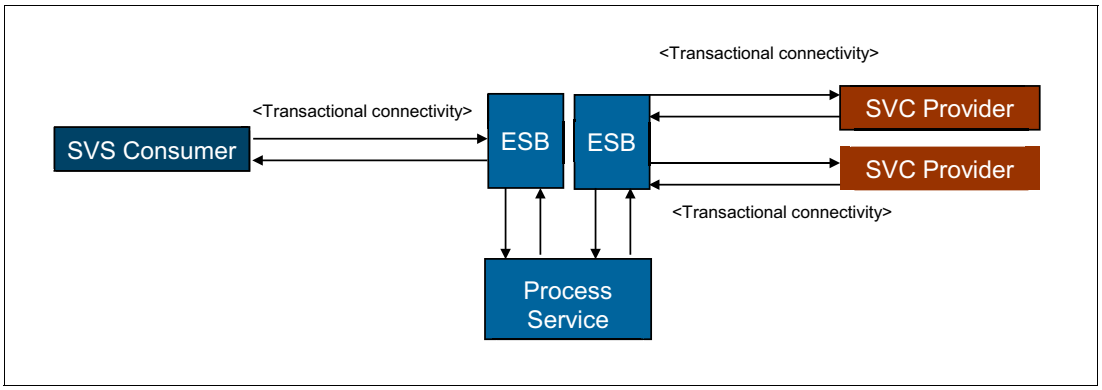

Figure 3-7 Integration pattern: Simple orchestration

Aggregation involves sending messages to multiple destination services, collecting the responses (perhaps asynchronous responses over a period of time), and combining them into a single response. *Sequential dependency* refers to scenarios where the overall business service involves interactions with multiple destination services but with dependencies between the calls.

In Figure 3-7, the ESB is used to provide mediation logic between the service consumer and the business service exposed by the process services component. The same ESB component provides mediation logic between the process services component and service providers. The process services component provides the logic required to start the service providers, handle the result messages and error situations, and formulate the response.

3.4 Architecture overview of ongoing integration with SAP

The reference architecture for the enterprise integration services components that support the integration patterns described previously is shown in Figure 3-8.

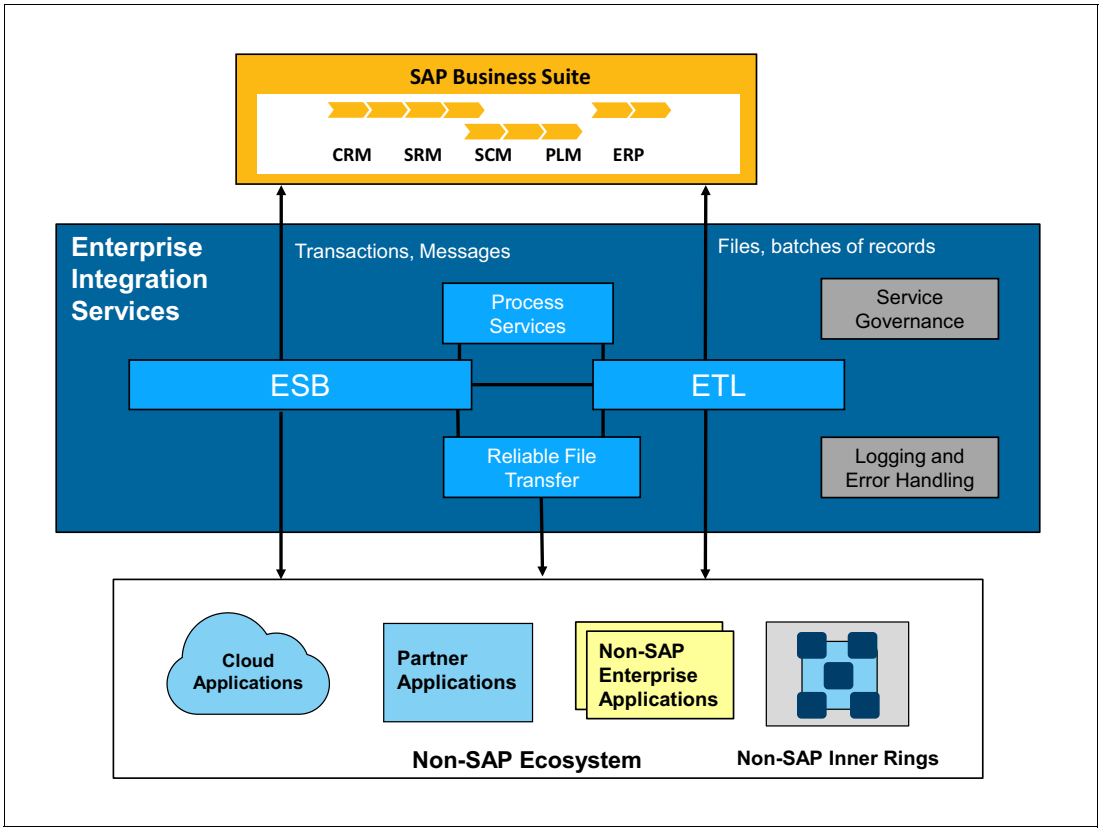

Figure 3-8 Enterprise integration reference architecture

Several components make up the enterprise integration services architecture:

- ESB
- ETL
- Service governance
- RFT
- Process services
- Logging and error handling

These components work together to provide the capabilities required to connect SAP systems with non-SAP applications within the enterprise, business partners, and cloud-based applications. In addition to business applications, existing non-SAP applications can include other ESB and ETL systems that already exist within the enterprise.

3.5 Architecture components of ongoing integration with SAP

This section introduces the key components of the architecture overview shown in Figure 3-8 on page 51, providing more detail and product mapping for the components.

3.5.1 Enterprise Service Bus

The ESB component is responsible for providing connectivity and integration logic for transactional interfaces. The primary function of the ESB is to decouple and isolate the application endpoints from one another, increasing the flexibility of the system and reducing the overall cost of integration.

The ESB provides this decoupling in several ways:

- Independent connectivity, with each application's endpoints using technologies that are appropriate for the functional and non-functional requirements of the interface, and using the available capabilities of the applications.
- Transformation and translation of messages from the source system to the destination system.
- Routing of messages. Routing includes simple routing across firewalls and other aspects of the infrastructure, in addition to more sophisticated business scenarios where rule-based selection is made between available endpoints.
- Ensuring required security is enforced using technologies and standards that address the interface requirements and the available capabilities of the application. This characteristic implies independent management of all of the required security artifacts, such as certificates.

The term ESB often implies the existence of reusable services, and represents a platform by which service consumers and service providers are connected. In this context, the ESB fulfills two core principles: Service virtualization and aspect-oriented connectivity.

Service virtualization

Service virtualization refers to the ability of the ESB to virtualize service interactions:

- Protocol and pattern. Interacting participants need not use the same communication protocol or interaction pattern. For example, a requester might require interaction through some inherently synchronous protocol, but the service provider might require interaction using an inherently one-way protocol with two correlated interactions. The ESB provides the conversion needed to mask the protocol and pattern switch.
- Interface. Service requesters and service providers need not agree on the interface for an interaction. For example, the requester might use one form of message to retrieve customer information, and the provider might use another form. The ESB provides the transformation needed to reconcile the differences.
- Identity. A participant in an interaction need not know the identity, for example, the network address, of other participants in the interaction. For example, service requesters need not be aware that a request can be serviced by any of several potential providers at different physical locations. The actual provider is known only to the ESB, and in fact, can change with no effect to the requester. The ESB provides the routing needed to hide identity.

Aspect-oriented connectivity

Aspect-oriented connectivity includes multiple cross-cutting aspects of integration, such as security, management, logging, and auditing. The ESB can implement or enforce

cross-cutting aspects on behalf of service requesters and service providers, removing such aspects from the concern of the requesters and providers themselves. Implementing these cross-cutting aspects within the middleware layer makes it easier to ensure consistent application of the logic, and reduces the amount of effort required to make changes.

In addition to the standard service-oriented definition of an ESB, the capabilities of the ESB component in this reference architecture have been somewhat extended to include more traditional message brokering (see Figure 3-9).

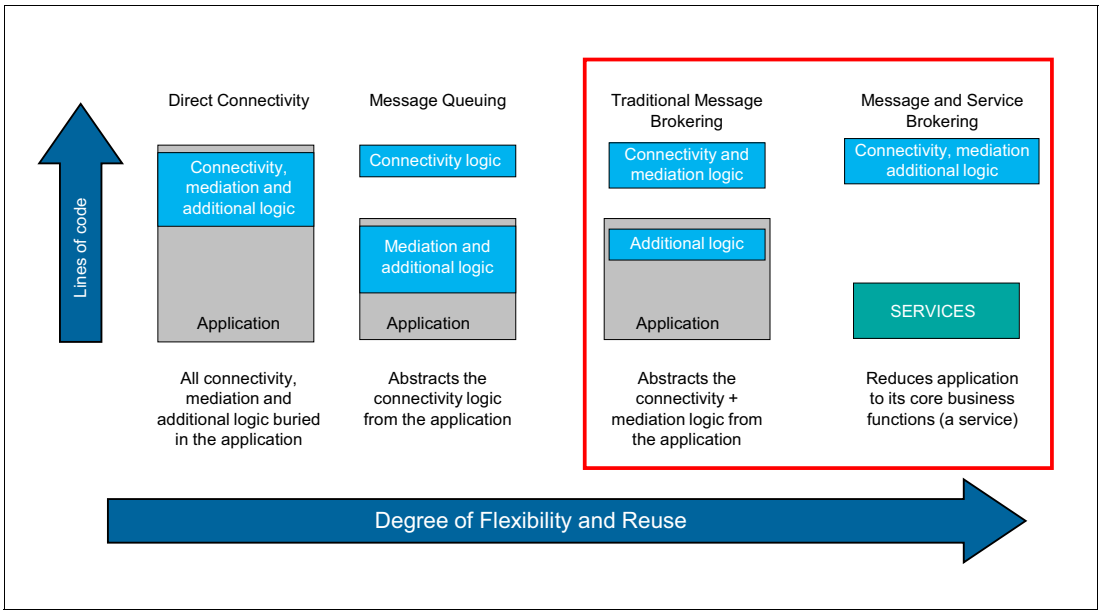

Figure 3-9 ESB functionality: Message and service brokering

Traditional message brokering is generally independent of the service reuse aspect of an ESB. In other words, the same or similar back-end service can be provided by different messaging endpoints in the message broker, which does not necessarily provide a single logical view of a service but still provides decoupling.

This decoupling occurs by removing connectivity and mediation logic from the individual application components and moving it into a shared middleware layer, the ESB. Moving the logic into the ESB layer creates opportunities for reusing the mediation logic between applications and interfaces, and it adds a degree of flexibility because the logic is centrally managed.

Identifying and implementing interactions between systems and services provides the greatest degree of flexibility and reuse. However, depending upon the long-term business value of a particular application or business function, it might not be cost-effective to implement the integration logic as a service. As a result, in a typical SAP deployment, the integration logic is created as a mix of traditional message brokering and service brokering.

Decoupling begins with the technologies used to send information across the network (the transports). The ESB supports a variety of transports and does not require that all participants in the interface communicate in the same way.

Typical transports used to interact with SAP systems include direct HTTP-to-enterprise services, posting flat files to the file system, and proprietary SAP technologies, such as RFC and IDocs. The IBM products providing ESB functionality support the proprietary SAP transports using IBM WebSphere Adapter for SAP Software.

The ESB also provides decoupling through message transformation. The ESB includes powerful tools for converting messages from one message format to another, both in terms of message structure and field semantics and formats.

Message transformation prevents the endpoints of the interface from needing to understand multiple message formats, or the details of remote systems with which it is interacting. This is particularly true in SAP deployments where much of the interaction involves proprietary BAPI and IDoc message structures, which are complex and difficult to understand.

IBM products for ESB

IBM has a range of products generally available for providing required ESB functionality. They are able to scale and perform as required by large SAP deployments.

IBM Integration Bus

IBM Integration Bus provides the principal integration platform for connecting SAP systems to other application components with real-time interactions. As shown in Figure 3-10, IBM Integration Bus supports transactional and batch-oriented integration patterns.

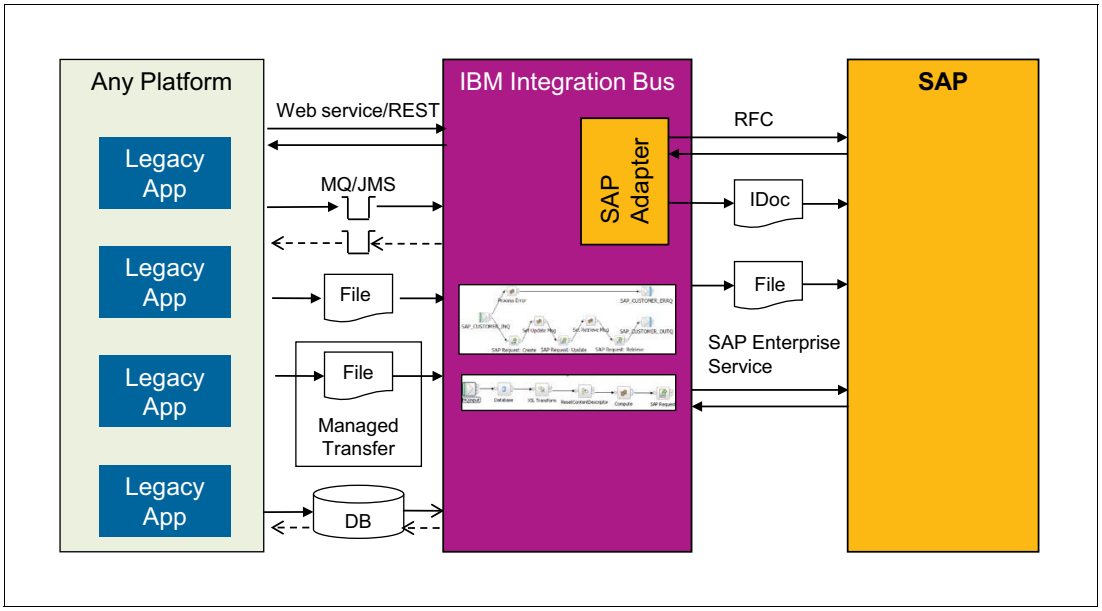

Figure 3-10 IBM Integration Bus: Overview

For batch-oriented integration, some limitations exist, which are described in 3.5.7, “Integration workload placement guidelines: ESB versus ETL” on page 64.

IBM Integration Bus supports both message and service brokering approaches, and provides integration through several transports and protocols:

- Integration with SAP BAPI, RFC, and IDoc through WebSphere Adapter for SAP Software
- Integration with SAP Enterprise Services using web services
- Integration with existing systems using web services, WebSphere MQ, the file system, and the database

IBM Integration Bus is pre-integrated with the IBM Security product suite to enable for authentication and authorization if needed, and identity mapping and identity propagation in a heterogeneous application landscape, including between non-SAP and SAP systems.

IBM Integration Bus is capable of handling thousands of messages per second, and performing complex mediation logic. In addition to those listed earlier, IBM Integration Bus supports a wide variety of other communication protocols, with several adapters available for other packaged applications. Common middleware functions, such as common logging, can be built as reusable sub-flows for use across all interfaces.

IBM WebSphere Transformation Extender

IBM WebSphere Transformation Extender provides advanced capabilities for mapping data from one message format to another. WebSphere Transformation Extender supports Extensible Markup Language (XML) and binary message formats with sophisticated tools to develop the mappings.

WebSphere Transformation Extender is not an ESB, but it integrates with IBM ESB technologies as an add-on.

Maps developed in the WebSphere Transformation Extender Design Studio can be used on multiple ESB platforms. They provide modular design, enabling projects to build a library of mapping logic that can be reused as-is, or composed into new transformation maps.

IBM WebSphere DataPower SOA Appliances

IBM WebSphere DataPower® is an appliance-based ESB that provides advanced security capabilities and XML processing at wire speeds. DataPower can be used as a gateway ESB to connect across security zones and between technical domains. Unfortunately, DataPower does not support the WebSphere Adapter for SAP Software, so it requires some other integration component, such as IBM Integration Bus, to integrate with SAP proprietary protocols using all of the integration patterns.

For several integration patterns, DataPower can provide value, and including it in the SAP integration logic is advantageous, using, for example, web services or WebSphere MQ. If DataPower exists within the environment, it can be added to the integration flow.

IBM WebSphere Cast Iron Cloud integration

IBM WebSphere Cast Iron® Cloud integration provides a cloud-based integration platform that can be deployed on-premises or as software as a service (SaaS). It supports a broad range of integration technologies, including proprietary SAP protocols, and a simple approach to developing integration logic. In most applications, IBM Integration Bus can address all of the integration requirements of an SAP project, but IBM WebSphere Cast Iron Cloud integration might be appropriate in specific situations.

3.5.2 Extract, transform, and load (ETL)

ETL is a term used broadly to refer to the activities required to move large volumes of data between systems in large batches. In the context of enterprise integration services, the ETL component is responsible for providing connectivity and integration logic for batch-oriented interfaces in ongoing integration (as opposed to initial data load or conversion activities). As the name suggests, an ETL flow has three major processes:

- Extract the data from the source system.
- Transform (and optionally cleanse) the data.
- Load the resulting data into the destination system.

ETL technologies are built to efficiently process very large sets of data, with internal staging of the data and parallel processing.

IBM products for ETL

This section introduces IBM products suitable for ETL.

IBM InfoSphere Information Server

Although the InfoSphere Information Server family includes several modules that support data integration activities, InfoSphere DataStage is the primary component used as the integration engine. It can transform and integrate any volume of information, and provides hundreds of built-in transformation functions.

DataStage provides direct connectivity with SAP systems through the InfoSphere Information Server Pack for SAP Applications. Three SAP interfaces are supported:

- BAPI. Extract and load
- IDoc. Extract and load
- ABAP. Code generated by the IBM Information Server Pack for SAP Applications to perform extract logic only

The ABAP Extract Stage shown in Figure 3-11 is just one out of many SAP interfaces supported by DataStage. The InfoSphere Information Server Pack for SAP Applications includes wizards that enable you to browse the SAP data dictionary and generate complete ETL jobs for IDoc, ABAP, and other interfaces with just a few mouse clicks.

This automated generation during the design drastically minimizes the effort compared to manual or template-based approaches, and it reduces errors while increasing quality and performance of ETL.

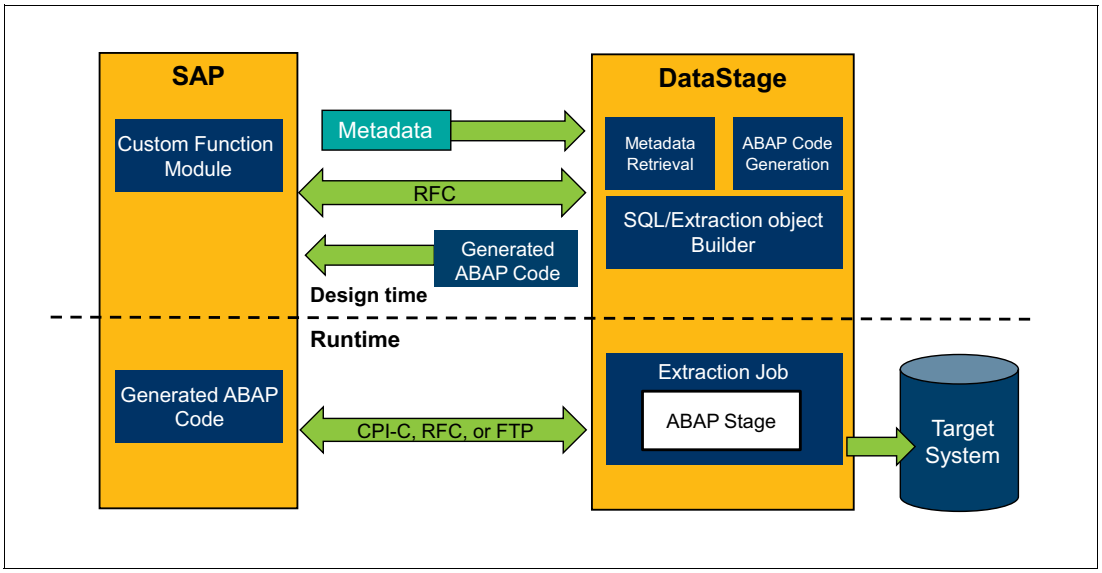

Figure 3-11 Solution architecture for ABAP integration with DataStage

The InfoSphere Information Server Pack for SAP BW provides connectivity for analytical SAP systems. For this purpose, it supports SAP interfaces such as Open Hub.

Additional information integration patterns, such as data replication and federation, can be implemented with InfoSphere Information Server. Data replication technology can be used to migrate an Oracle database supporting an SAP application to IBM DB2.

Other important characteristics of InfoSphere Information Server are the broad information governance capabilities that are required to manage critical aspects of information assets, such as information quality, information lifecycle, information protection, and so on. These capabilities of IBM InfoSphere Information Server can be used to manage, for example, information quality for SAP applications, optimizing business processes with high-quality data.

These functions are also used by the Master Data Management (MDM) component of the IBM Reference Architecture for SAP described in Chapter 7, “Master data management for SAP” on page 159.

IBM Integration Bus

Although IBM Integration Bus is typically used in transactional integration, it is also a good platform for handling simple ETL interactions. IBM Integration Bus supports batch ETL transports for both SAP and non-SAP systems (with the exception of the ABAP Stage functionality available in InfoSphere Information Server).

IBM Integration Bus can handle moderately sized messages directly, and large messages by splitting along record boundaries (assuming the message is made up of multiple records). IBM Integration Bus can also be used in scenarios where internal staging is required, but in these scenarios it requires more development effort than DataStage.

In general, IBM Integration Bus can be used extensively to handle simple ETL interfaces. More detailed guidelines regarding workload placement in IBM middleware products is provided in 3.5.7, “Integration workload placement guidelines: ESB versus ETL” on page 64.

3.5.3 Service governance

Service governance includes, among others, two important aspects:

- Service lifecycle management
- Service run time

The following sections provide more details about each.

Service lifecycle management

The service lifecycle management aspect of service governance provides a set of processes implemented to ensure proper sharing and reuse of services throughout the service lifecycle. It is often described in terms of the following activities:

- Publish and find. Make services available by publishing them to a central repository. Search for existing services for reuse.
- Enrich. Integrate with runtime environments and govern service design and runtime policies.
- Manage. Monitor and report on service usage. Manage service contracts and multiple service versions.
- Govern. Maintain and govern service contracts, service consumers, and service providers. Integrate with other products that support service-oriented architecture (SOA) governance.

In practical terms, the service governance component provides a service repository for storing service artifacts, and a customizable, ready-to-use process for managing those service artifacts through the project lifecycle.

The following scenario provides an example:

1. New integration requirements are identified. Using the information provided, for example business data and business operation, the integration architect or developer searches the service repository for existing services:
   - If a matching service is found, the service is evaluated to determine if enhancements are required to support the new requirements.
   - If a matching service is not found, define a new service and publish it to the repository.
2. In addition to business data and business operation requirements, there are also technical requirements for the service. They are identified and implemented as part of the middleware implementation of the service, including attaching appropriate policy entries.
3. When the service is ready, promote the service to the next state in the service governance lifecycle. At appropriate times in the process, review and approve the defined services. As services are promoted through the stages, the service governance component pushes the changes out to the relevant runtime instances.
4. As new service consumers are added, define the appropriate entries in the repository for consumers and their related service level agreements (SLAs) and policies.
5. Retrieve monitoring information from the runtime environment and validate the actual service consumption.

Service run time

In addition to design-time processes, service governance includes a runtime service registry where the defined service policies and configurations are enforced by the service integration components. The resulting runtime information is provided to report on service use.

In practical terms, the ESB and ETL components connect to the service governance registry to get information about the configured services, and ensure that the policies are applied to the messages that pass through the system. These enforced policies include consumers that are authorized to use a particular service, routing policies, endpoint references, and transformation policies.

The ESB solution can be implemented in conjunction with the service registry to provide configuration-driven integration. In this approach, the ESB is implemented with a common flow for each integration pattern addressing common mediation functionality that is driven by external configuration information.

A simple example message flow from IBM Integration Bus is provided in Figure 3-12.

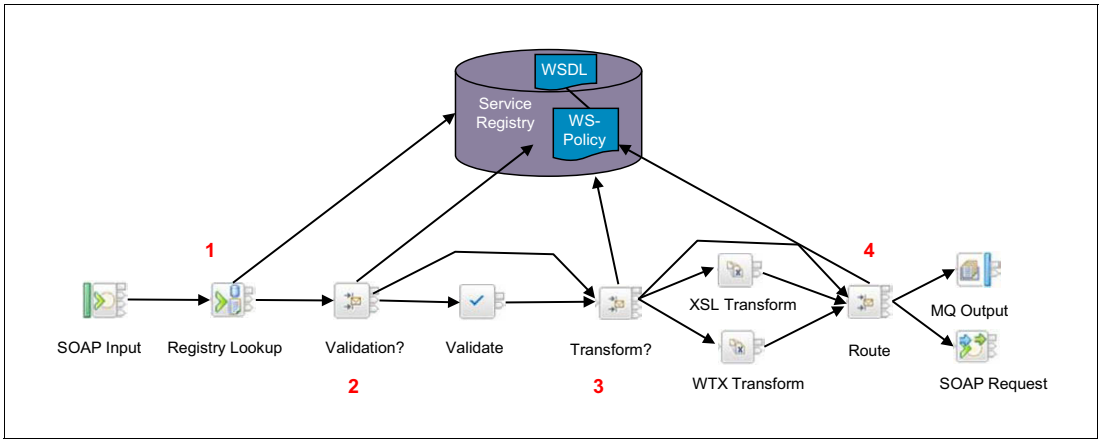

Figure 3-12 ESB flow using WebSphere Service Registry and Repository (WSRR) policy configuration

The following steps are shown in Figure 3-12 on page 58:

1. At run time, when the ESB receives a service request, it uses information in the request to look up the service definition in the service registry.

   The result from the service registry contains the service definition and the associated policy definitions. The policy definitions can include any of the following items:

   – A requirement for validation and the schema that should be used
   – The transformation logic required, such as a reference to the Extensible Stylesheet Language Transformation (XSLT) file or WebSphere Transformation Extender (WTX) map
   – The routing logic that should be applied
   – The endpoints for service providers
2. If validation is required, as defined by the service configuration, the ESB applies the provided schema.
3. If transformation is required, the ESB retrieves the transformation logic from the service definition and applies it to the incoming message.
4. The endpoint routing is determined by the endpoints defined in the service registry using the defined transport protocol.

Because the mediation logic is defined in the external service registry, the ESB flow can support multiple service implementations without the need to develop ESB-specific logic. If changes are required, they can be applied easily in the service registry, without requiring code deployment to the ESB, which is usually more complicated.

Finally, because usually each service also exposes data to the service consumers, in mature SOA environments a consideration in the service design is to check whether it is appropriate, through the service, to expose the data outside the service. The reason for this check is simple.

Low quality data, or, generally speaking, data that is not fit for use outside the component that holds it, can be easily exposed through a service, but behaves like a virus in the enterprise. It “infects” other systems that have good data, and are now consuming poor data through a service.

In mature SOA environments, service governance intersects with information governance. Applying information governance includes the use of tools to measure and control all dimensions of data, such as quality, lifecycle, security, and so on. Enforcing standards around the data aspect of service design is supported through the information governance infrastructure.

IBM products for service governance

IBM WebSphere Service Registry and Repository (WSRR) is the master metadata repository for service definitions and related service metadata. It applies to both traditional web services that implement WSDL interfaces with SOAP/HTTP bindings and to a broad range of SOA services that can be described using WSDL, XML Schema Definition Language (XSD), and policy decorations. WSRR can use a range of protocols, and be implemented according to a variety of programming models.

As the integration point for service metadata, WSRR establishes a central point for finding and managing service metadata. The service metadata is acquired from several sources:

- Service application deployments
- Other service metadata
- Endpoint registries
- Repositories, such as Universal Description, Discovery, and Integration (UDDI)

WSRR is where service metadata that is scattered across an enterprise is brought together to provide a single, comprehensive description of a service. Through this approach, visibility is controlled, versions are managed, proposed changes are analyzed and communicated, usage is monitored, and other parts of the SOA foundation can access service metadata with the confidence that they have found the copy of record.

Specifically for SAP deployments, WSRR can be integrated with the SAP Enterprise Service Registry to extract Enterprise Service definitions and make them available across the enterprise.

ESB technologies from IBM, such as IBM Integration Bus, provide tight integration with WSRR. This integration enables the ESB component to use the service configuration defined with WSRR and enforce defined policies, such as SLAs with service consumers and applicable routing and transformation logic.

3.5.4 Reliable File Transfer

Solutions that use traditional file transfer technologies suffer from a common set of challenges:

- Management of the connectivity is performed in a distributed fashion, and handled separately for each system.
- Error handling must be managed by the endpoint applications, for example, restarting a failed file transfer.
- Monitoring of transactions is limited.
- Security is inconsistent and complex.

Figure 3-13 shows a tightly coupled and brittle environment that results from solutions that use traditional file transfer technologies.

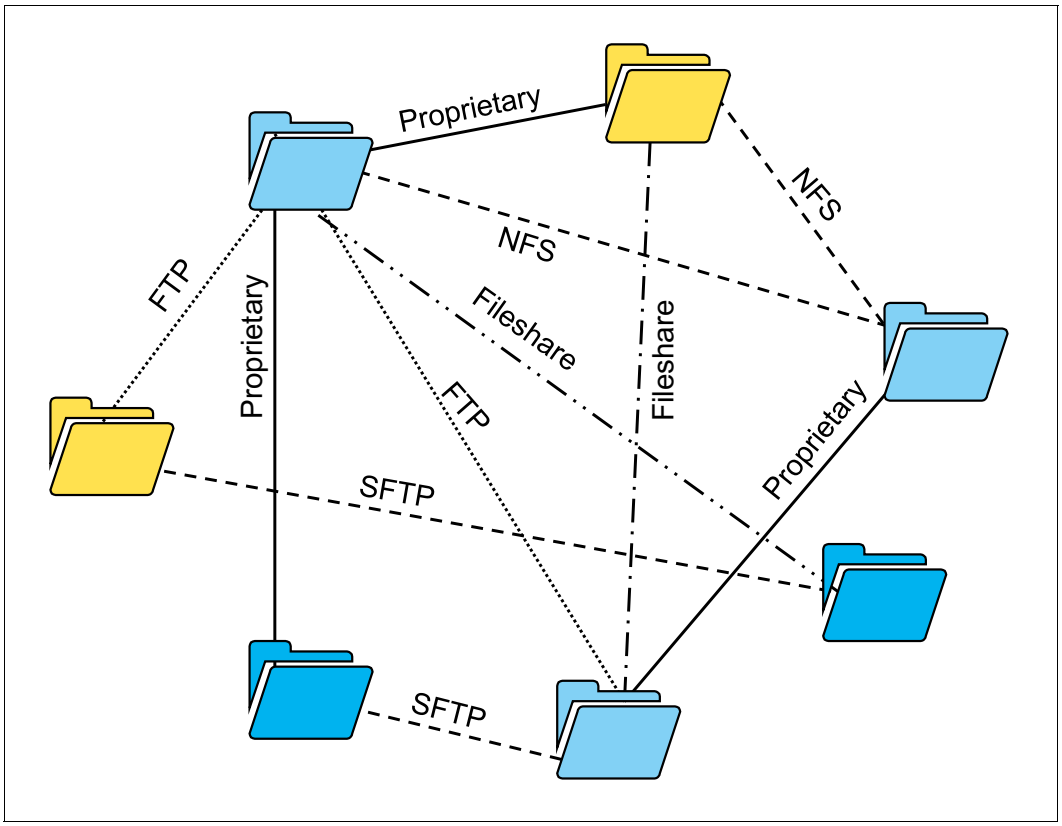

Figure 3-13 Example scenario with traditional file transfer technologies

Figure 3-14 shows an example of an environment that takes advantage of Reliable File Transfer technologies.

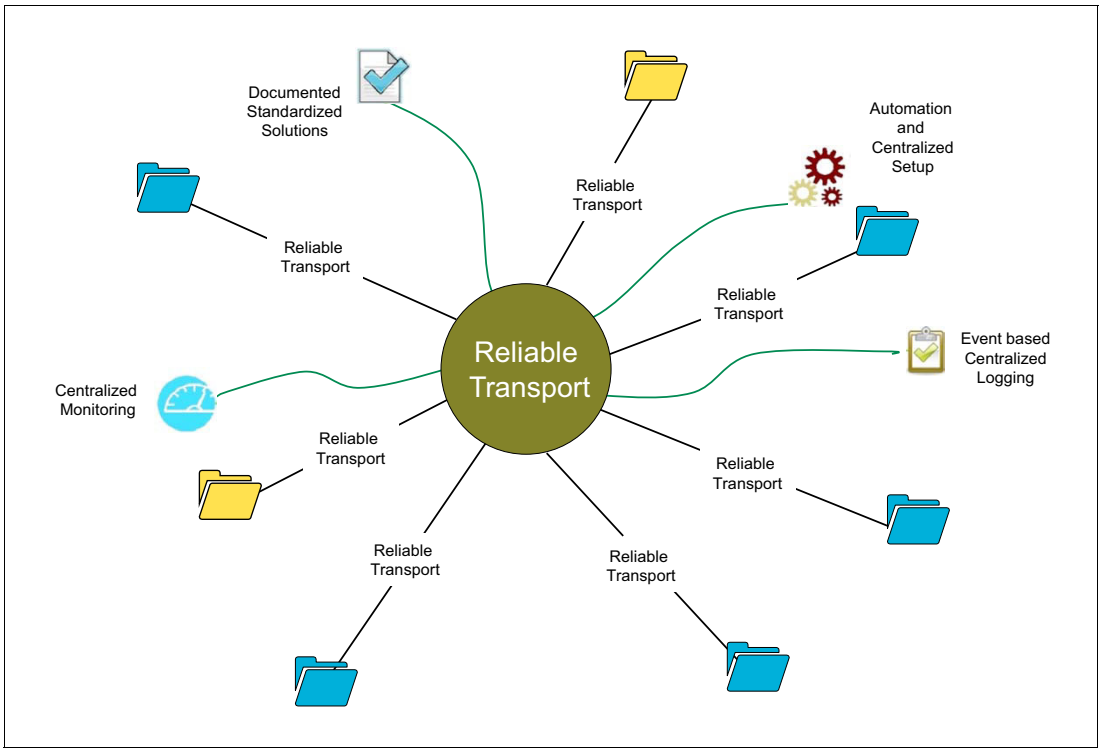

Figure 3-14 Example scenario that takes advantage of Reliable File Transfer technologies

An environment such as the one shown in Figure 3-14 provides central configuration and setup, centralized logging and monitoring of file transfers, and a standard solution with established quality of service characteristics for implementing file transfers within the enterprise. Additionally, because the configuration for the Reliable File Transfer environment is centrally managed, the resulting solution provides loose coupling and flexibility.

In a typical SAP implementation, several scenarios require the movement of files from one point to another. Most of these situations involve existing non-SAP application components that have been in operation for a long time, where file-based interactions are the most convenient, and in some cases only, way to move the data.

Typically, the middleware integration components (ESB/ETL) are the recipients of the files. With Reliable File Transfer technologies, the underlying transport can often be changed with little effect to the application endpoints.

IBM products for Reliable File Transfer

This section introduces IBM products suitable for solutions based on Reliable File Transfer technologies.

IBM WebSphere MQ File Transfer Edition

IBM WebSphere MQ File Transfer Edition enables organizations to manage file transfers in WebSphere MQ environments across a range of platforms and networks. WebSphere MQ File Transfer Edition builds upon existing WebSphere MQ networks, which are ubiquitous in many enterprises. Through WebSphere MQ File Transfer Edition, WebSphere MQ queues are effectively extended into the file system, and can be managed in a central place.

The following list includes some features and characteristics of WebSphere MQ File Transfer Edition:

- Provides a customized, scalable, and automated solution, enabling managed, trusted, and secure file transfers while eliminating costly redundancies
- Uses existing messaging infrastructure for universal service delivery, including messages, files, and events
- Provides end-to-end audit trail across file transfers
- Facilitates a secure and reliable managed file transfer environment across IBM Sterling Connect:Direct® and WebSphere MQ File Transfer Edition endpoints
- Provides reliable transfer of file data between internal systems and B2B gateways
- Enables applications (for example, office productivity tools and computer-aided design (CAD) programs) to use web application programming interfaces (APIs) to move files
- Modernizes batch-oriented architecture into micro-batches with simple messaging conversion

IBM Sterling Connect:Direct

IBM Sterling Connect:Direct is the point-to-point file transfer software optimized for high-volume, secure, assured delivery of files within and among enterprises. IBM Sterling Connect:Direct can deliver files with the following characteristics:

- Predictability. Assures delivery using automated scheduling, checkpoint restart, and automatic recovery and retry.
- Security. Ensures that the information of an organization stays private, and that the file transfers are auditable for regulatory compliance using a proprietary protocol, authorization, and encryption (Federal Information Processing Standard (FIPS) 140-2, and Common Criteria certified).
- Performance. Handles the most demanding loads of an organization, from high volumes of small files to multi-gigabyte files.

IBM Sterling Connect:Direct is available as licensed on-premises software.

3.5.5 Process services

Process services are technical integration processes that provide advanced integration logic beyond the typical mediation that is provided by ESB components.

Although process services are typically implemented on a business process management (BPM) platform, a clear delineation should be drawn between business processes that provide the logic for business operations (including human interaction by business users) and process services that provide support for technical integration requirements. Even when process services can involve human interaction in some cases, the users involved in those activities are typically in technical support roles.

The following list includes some of the use cases in which process services are used:

- Service aggregation
- Correlating (asynchronous) requests and responses
- Selecting responses from multiple service providers (1 to N)
- Orchestrating sequential execution of multiple, dependent services
- Transaction management
- Complex error handling

IBM Products for process services

This section introduces IBM products suitable to implement process services.

IBM Integration Bus

IBM Integration Bus can be used to implement simple process services.

IBM Business Process Manager Advanced

IBM Business Process Manager Advanced process server enables the deployment of standards-based integration logic in an SOA, which provides and consumes services in a defined executable business process. As a process service engine, the use of IBM Business Process Manager Advanced process server is restricted to integration components that provide automated technical services.

3.5.6 Logging and error handling

The logging and error handling component is responsible for recording and tracking application-level errors, and for providing functions to address those errors. This component work in conjunction with infrastructure-level monitoring to provide a comprehensive view of system operations.

All of the components in the enterprise integration services subsystem send events that are received by the error management component and recorded. After the events are received, the logging and error handling component can perform the following tasks:

- Display the event information.
- Send notifications to the appropriate personnel.
- Perform automated corrective actions.

In some cases, the corrective action can be to resubmit the original message back through the middleware components. However, care must be taken when using this approach to ensure that the resubmitted failed transactions do not overwrite subsequent successful transactions. This situation can be addressed by designing the interface messages with information required to detect stale data, or by limiting the scenarios for which the resubmission is enabled (for example, batch processes that are run infrequently).

3.5.7 Integration workload placement guidelines: ESB versus ETL

For the most part, the individual components within the enterprise integration services have distinct capabilities, and do not overlap in their functionality. One exception is those interfaces that involve batch processing, because both the ESB and ETL components have capabilities that are able to satisfy many batch integration scenarios.

For complex batch interfaces that require staging or cleansing, or that might require parallel processing for very large amounts of data, the ETL component is best equipped to satisfy the requirements.

The following questions help to determine which component is best suited to address a particular scenario:

- What is the nature of the integration? Is it ETL and data synchronization, or transactional?

  In some cases, the line between ETL and transactional integrations is a thin one. Generally, transactional integrations encompass single query, create, or update transactions, and ETL integrations deal with loading large amounts of data.

- Does an ETL job already exist from initial data load that can be reused?

  In many cases, the initial data load is performed to move data from an existing system into a new system, where the new system assumes the business function and the existing system will be withdrawn from service. However, in other scenarios, the existing system remains in place, and the information must be synchronized both as an initial data load and an ongoing interface. In those situations, the interface design and related integration logic for the initial data load and ongoing efforts should be harmonized.

  The assumption is that the ETL components are used to implement the initial data load logic from existing applications into the SAP environment. This initial data load activity should cover most, if not all, of the data elements that are required later for incremental data synchronization. If this is the case, there is an opportunity to reuse the existing logic.

  A key consideration is not just the existence of the data job, but also the level of reuse that can be achieved. In many cases, a data job created for a one-time, initial data load might not be robust or complete enough for the purpose of ongoing data synchronization without additional development. The relationship between the initial data load and the ongoing interface must be taken into consideration during the design phase.

- Will the ETL job provide sufficient decoupling between the integration points?

  One of the fundamental principles of SOA (and of integration architecture in general) is decoupling the systems involved in the interaction from one another. Decoupling pertains to several aspects of the integration between the source and destination systems, including but not limited to the following aspects:

  - Transport technologies between source and destination, and between multiple sources
  - Message formats
  - Scheduling and frequency of the interface

- Does the data integration require complex transformation logic?

  Often, ETL jobs require complex transformation logic on large volumes of data. InfoSphere DataStage provides facilities for parallel processing of complex data transformations to achieve high throughputs.

- Is data cleansing required before the data is delivered to the destination?

  Occasionally, the data provided for ETL requires cleansing and data matching. Traditional ESB technology does not support this functionality easily, but this functionality is a core capability of DataStage and IBM InfoSphere QualityStage®.

- What are the non-functional requirements for the integration, for example, number of records per message, size per record, frequency of integrations, and so on?

  IBM Integration Bus has limits on the size of messages that can be reasonably processed. A well-established design pattern is to break large messages into smaller batches for processing. The contents of the large message are read into memory when the smaller batch of records is loaded, limiting the memory usage for the message flow. In this case, IBM Integration Bus defines a large message as a size between 5 MB and 100 MB.

  Besides memory usage, processor usage must also be considered when handling large messages in IBM Integration Bus. The performance characteristics of IBM Integration Bus with regard to large message processing have been tested and documented in the IBM Integration Bus Performance Report. The characteristics vary somewhat based upon

individual record size and the message format (XML versus delimited or fixed-length format). However, as is to be expected, larger messages consume more resources, take longer to process, and enable a lower throughput than smaller messages.

As an example, if the message flow involves taking in a file and writing out a file, where the incoming message format is XML, an 8 MB message can be processed at 0.46 messages per second consuming 4467 processor milliseconds (ms) per message.

Performance improves with delimited message formats, processing 3.38 messages per second and consuming 454 processor ms per message. Although performance figures are not available for DataStage at this point, processing large messages is a common scenario for the platform.

The DataStage architecture establishes a grid of sorts to process the large messages in parallel, and deal with the results simultaneously. Additionally, using DataStage for large message processing provides a partitioning of the workload that keeps the IBM Integration Bus resources freed up to handle more transactional workloads.

- What are the transactional and error handling requirements of the integration?

  The drawback of handling large messages within IBM Integration Bus is the transactional characteristics of the message flow. If a large message is received, it is processed record by record, and failed records are handled individually also. Depending upon the requirements for the interface, this behavior might be wanted. However, in other cases, an all-or-nothing approach to processing the message is more appropriate.

3.6 Initial data load

The two common scenarios for initial data load in SAP solutions are as follows:

- Deployment of a new SAP application instance. This use case primarily occurs when a company (often times after a series of mergers and acquisitions) decides to consolidate a heterogeneous application landscape by replacing existing applications with SAP applications. This use case requires massive data harmonization from various sources before it can be loaded to SAP systems.
- Consolidation of SAP applications. Some enterprises implemented SAP R/3 by country or line of business (LOB), and ended up with dozens or hundreds of SAP application instances. In such scenarios, often as part of upgrading to the SAP Business Application suite, the SAP application instances have been consolidated to fewer instances, or just one instance, for a particular application, such as SAP ERP.

  Because no two SAP instances are configured in the same way, this consolidation requires massive data harmonization toward the new application instance.

The SAP AcceleratedSAP (ASAP) methodology for implementation is the SAP roadmap for implementing SAP solutions in a cost-effective, speedy manner. In the SAP ASAP method, the process of data harmonization to get the data load-ready is known as *data conversions*, and the specifications are the CDD documents mentioned in 3.2.1, “Align enterprise integration services with SAP implementation methodology” on page 41.

IBM offers a unique solution for data conversion, known as IBM InfoSphere Information Server Ready to Launch for SAP Applications, which provides an industrialized approach to migrate existing data to SAP systems. This solution is composed of three major components:

- A proven delivery methodology aligned with the SAP ASAP methodology, reducing the SAP implementation risk.
- IBM InfoSphere Conversion Workbench for SAP Applications. This product provides a unique set of capabilities targeted toward data migration for SAP applications, drastically reducing time and efforts while improving data quality with a superior degree of automation. InfoSphere Conversion Workbench for SAP Applications is built on InfoSphere Information Server.
- InfoSphere Information Server. This product is the enterprise information integration platform from IBM.

This solution has the following major benefits:

- Reduction of time and deployment efforts through a high degree of automation of previously manual tasks, which also reduces errors
- Risk mitigation through proven methodology
- Improved business process execution through high-quality information
- Delivery of a repeatable and reusable infrastructure that can be used for multiple SAP system rollouts, and for ongoing enterprise integration, quality, and governance

A unique differentiator compared to template-based approaches available in the marketplace is that, if the SAP target application changes, (for example new Z-tables, new Z-fields, changes in the SAP check tables storing reference values such as country codes, and so on), InfoSphere Information Server Ready to Launch for SAP Applications is capable of detecting such changes. Rather than manually adjusting the ETL logic, the ETL logic can be regenerated to reflect the changes.

For large SAP implementations with 60 - 80 SAP business objects in scope, each composed of several data tables with several dozen related check tables, such changes of the SAP target application during SAP Blueprint and SAP Realization phase occur frequently. The effort to adjust the templates is substantial, but with InfoSphere Information Server Ready to Launch for SAP Applications approach it is easy.

Another unique differentiator of the InfoSphere Information Server Ready to Launch for SAP Applications solution is that it brings data into focus early in the SAP Blueprint phase, with features like Business Data Roadmap (BDR). This is important because traditional approaches to data migration start to look at data much later, as part of the SAP Realization phase. When the SAP Realization phase starts, project timeline, budget, and so on, are already established.

If, after source system analysis, the data quality is worse than anticipated, it can cause project delays and budget overruns, because the load-ready data is in the critical path of going live. The InfoSphere Information Server Ready to Launch for SAP Applications solution takes data off the critical path by looking at data early in the SAP Blueprint phase using methodology and tools.

Figure 3-15 shows a conceptual overview of the InfoSphere Information Server Ready to Launch for SAP Applications solution. A complete description of this solution is beyond the scope of this book.

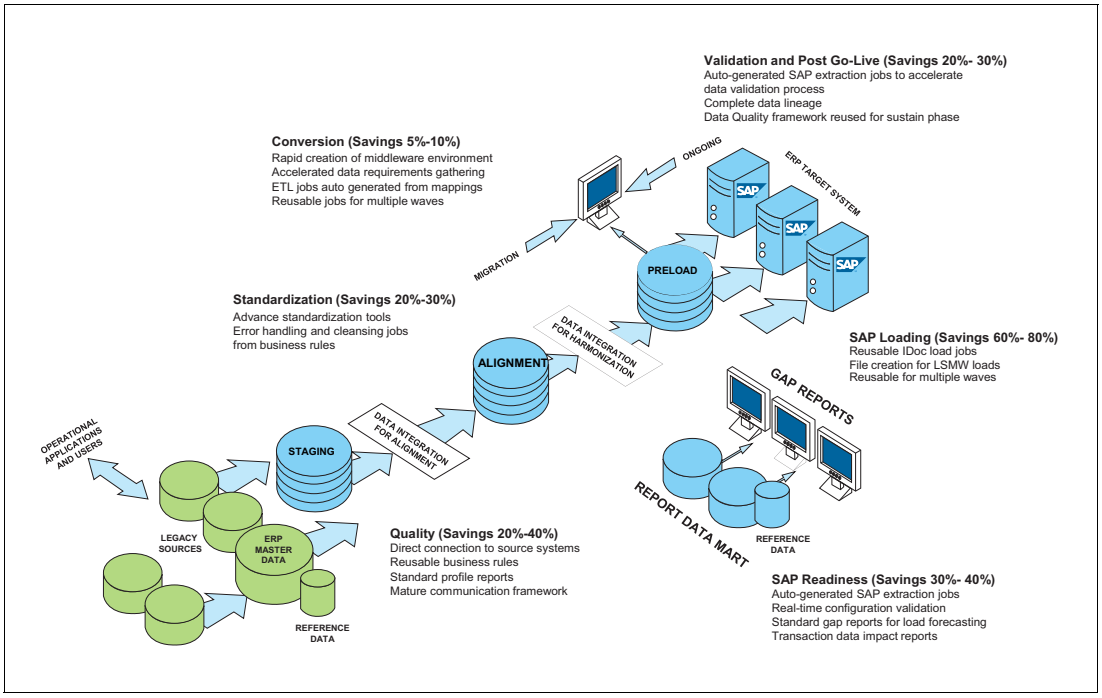

Figure 3-15 IBM InfoSphere Information Server Ready to Launch for SAP Applications: Overview

For the SAP Blueprint phase, the IBM InfoSphere Information Server Ready to Launch for SAP Applications solution provides these items:

- Business Data Roadmap

  This is a capability of the InfoSphere Conversion Workbench for SAP Applications. It enables the functional data analyst to capture all of the functional data requirements of the SAP target systems. The requirements can be captured early by participating in the process workshops during the SAP Blueprint phase.

  The functional data analyst does not come “empty-handed” to the process workshops. With the help of BDR, the functional data analyst can import the SAP business process hierarchies, associated business objects, and the data table structures that are associated with the business objects. The functional data analyst can then seamlessly capture the attributes that are in or out of scope, using the BDR web-based user interface (UI).

  Some business objects, for example, master data, appear in multiple process domains such as opportunity-to-order (OTO), order-to-cash (OTC), and so on, which are handled in separate process workshops. The functional data analyst can seamlessly see conflicting data definitions across process domains with the help of the BDR.

Figure 3-16 shows a sample page of the BDR web-based UI.

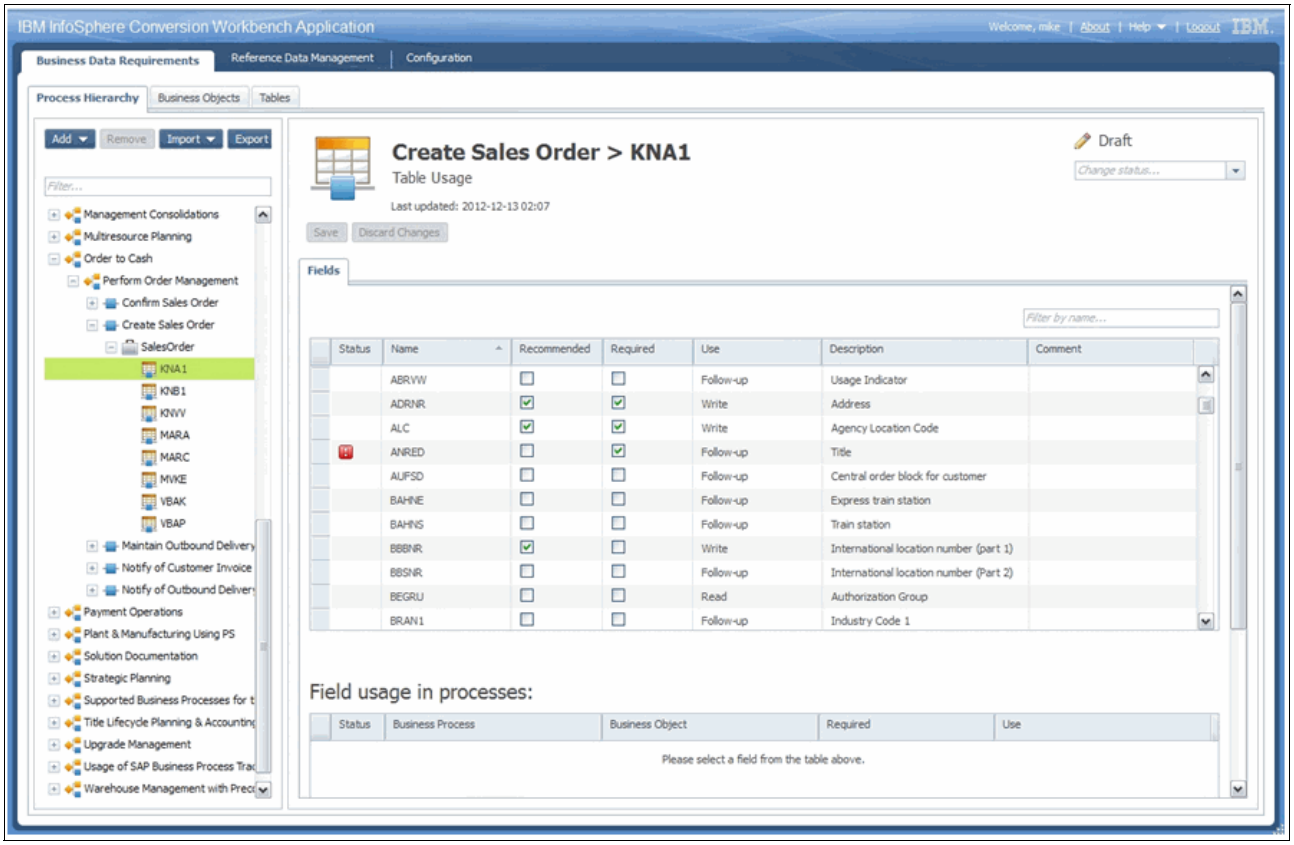

Figure 3-16 IBM InfoSphere Conversion Workbench for SAP Applications: BDR web-based UI

- Staging area

  While the data is moved from source to target, staging areas exist where the data is persisted while in transit. The staging area is modeled identically to the source system data models.

  The following list describes the key design points for the staging area:

  – The need exists to have a place to store the source data to avoid having to extract the data every time a test must be performed during the SAP Realization phase. Extracting the data every time has a negative effect on the source systems, which are still production systems at this point.

  – Data profiling is an intensive input/output (I/O) task that can have a negative effect on the performance of the source systems. IBM InfoSphere Information Server Ready to Launch for SAP Applications includes capabilities, such as Rapid Modeler and Rapid Generator, that can be used to extract data from existing SAP systems with just a couple of mouse clicks. The extracted data is moved into the corresponding staging areas.

    For non-SAP systems, the InfoSphere Information Server platform provides suitable capabilities. InfoSphere Information Server provides data profiling capabilities that help to analyze the data quality issues of the data sources in the staging area for the source.

Because the SAP target specification has been identified with the help of the BDR, performing a fit-gap analysis at this point between source models and source data quality, and the SAP target, results in the Data Quality Action Plan (DQAP). The DQAP defines the needed data cleansing activities for the SAP Realization phase. The logical source to target mappings are also defined during this phase.

For the SAP Realization phase, the IBM InfoSphere Information Server Ready to Launch for SAP Applications solution provides several key features:

- Architecture for alignment and preload areas (Figure 3-15 on page 68)

  Both areas are modeled by extracting the SAP target data model for the required scope from the SAP target system using IBM InfoSphere Rapid Modeler for SAP Applications. The InfoSphere information architect might remove for the alignment area a couple of the constraints from the SAP target model to enable all data that has not been cleansed from the various sources to enter the alignment process in the alignment area.

  The preload area is a one-to-one representation of the SAP target model. Therefore, only records that are compliant with the SAP target model can pass from the alignment area to the preload area.

  The following list describes the key reasons for this architectural design:

  – From the various staging areas (one per source) to the alignment area, structural alignment to a single common model is done by implementing InfoSphere Information Server-appropriate data model conversions.

    In a second step, semantic alignment, called *transcoding*, is performed to replace the various reference data values from the various sources with the corresponding reference data values from the SAP target system.

    When the data is structurally and semantically aligned in the alignment area, a single and common set of cleansing routines can be applied to all of the data records across all of the sources.

  – Cleansing tasks, such as name and address standardization, matching and deduplication, assignment of default values for mandatory fields in the SAP target system where sources did not have values, and so on, are completed using IBM InfoSphere Information Server Ready to Launch for SAP Applications. After cleansing tasks are complete, the data is transformed to a preload model. From the preload area, all data can then be loaded into the SAP target system.

- Reference data management for transcoding

  With IBM InfoSphere Conversion Workbench for SAP Applications (part of IBM InfoSphere Information Server Ready to Launch for SAP Applications), you can automatically download, from sources and the SAP target, the reference data tables into the IBM InfoSphere MDM Reference Data Management Hub application, which means you can more efficiently manage reference data.

  The functional data analyst uses the IBM InfoSphere Information Server Ready to Launch for SAP Applications web UI to define transcoding tables by aligning the source reference data values with their appropriate SAP target reference.

  After the functional data analyst defines the transcoding tables using this capability, the transcoding tables are pushed by the IBM InfoSphere Conversion Workbench for SAP Applications into the ETL environment. ETL routines use the transcoding tables as part of the semantic alignment to replace the source reference data values with their SAP target reference data values.

- Gap reports

  The gap reports measure key data quality characteristics to determine load readiness of the data for the SAP target system. The gap reports can run on the alignment and preload area. The gap reports can measure various metrics:

  – Completeness.

  – Category completeness. For example, customer records in different account groups can have different completeness requirements.

  – Validity. Compliance with reference data values in the check tables of the SAP target system, cross-business object dependencies, for example, the order object depends on customer and material, and so on.

  The beauty of the gap reports is that they are driven by the SAP metadata and configuration of the SAP target system. Therefore, the measurement logic is dynamically generated based on the current state of the SAP target system just before execution. Therefore, it truly measures load readiness from the SAP target system perspective.

  The following list describe the benefits of the gap reports:

  – Enable project management to see how much progress, in terms of correcting data quality issues, has been made since the gap reports were previously run. Gap reports can run daily, weekly, and so on, depending on project needs.

  – Enable you to determine whether the data is load-ready, because they measure the constraints enforced by the SAP interfaces.

  – Gap reports measure *actual data quality constraints*, because the measurement logic is generated just before running the gap reports. Manual errors are avoided because of generation of the measurement logic.

  – Enable you to view (as shown in Figure 3-17 on page 72) the data quality issues by data quality exception type per field, record, table, and SAP business object level, with appropriate drill-through capabilities.

Figure 3-17 shows sample gap reports in IBM InfoSphere Information Server Ready to Launch for SAP Applications.

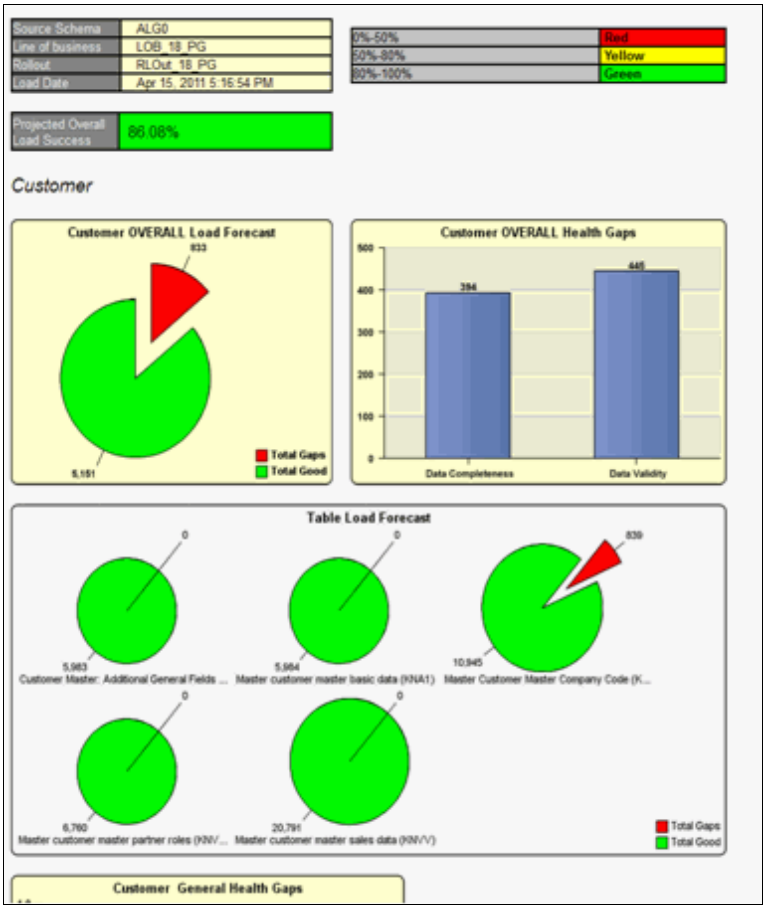

Figure 3-17 Sample gap reports

3.7 References

These websites are also relevant as further information sources:

- IBM InfoSphere Information Server Pack for SAP BW

  http://www.ibm.com/software/products/en/infosphere-information-server-pack-sap-bw
- IBM InfoSphere Information Server Pack for SAP Applications

  http://www.ibm.com/software/products/en/infosphere-information-server-pack-sap-applications
- IBM MQ

  http://www.ibm.com/software/websphere/ibm-mq.html
- IBM InfoSphere Information Server Family

  http://www.ibm.com/software/data/integration/info_server/
- IBM WebSphere Adapter for SAP Software

  http://www.ibm.com/software/products/en/websphere-adapter-mysap
- WebSphere Adapter for SAP Software documentation

  http://ibm.co/1lNkkxi
- IBM WebSphere Transformation Extender

  http://www.ibm.com/software/products/en/wdatastagetx
- IBM WebSphere DataPower SOA Appliances

  http://www.ibm.com/software/products/en/datapower
- IBM WebSphere Cast Iron Cloud integration

  http://www.ibm.com/software/products/en/castiron-cloud-integration
- IBM InfoSphere DataStage

  http://www.ibm.com/software/products/en/ibminfodata

4

Process optimization for SAP

IBM Smarter Process is a set of business process enhancement and optimization tools that work together to improve the business outcomes of business processes. IBM software products for Smarter Process include IBM Business Process Manager, IBM Operational Decision Manager, IBM Business Monitor, and IBM WebSphere Business Events.

IBM Smarter Process for SAP is a set of integrated capabilities contained in the IBM Smarter Process software product offerings, listed previously, that can improve the delivery of SAP implementations, and amplify SAP business value after the implementation.

Many of these capabilities are available in every edition of IBM Business Process Manager (Express, Standard, and Advanced), although comprehensive technical integration using SAP technical services requires IBM Business Process Manager Advanced and the IBM WebSphere Adapter for SAP Software.

This chapter explores how IBM Smarter Process for SAP capabilities can be used to accomplish the following goals:

- Decrease SAP implementation complexity, and lower cost and risk.
- Improve strategic alignment of the SAP solution.
- Improve the visibility, flexibility, agility, and control of SAP processes as they are running.
- Provide an effective platform for continuous active business performance optimization.

This chapter includes the following topics:

- 4.1, “SAP solutions as a system of engagement” on page 76
- 4.2, “Architecture overview” on page 77
- 4.3, “SAP active business performance optimization architecture” on page 78
- 4.4, “IBM Smarter Process for SAP capabilities” on page 81
- 4.5, “IBM Smarter Process for SAP products and solutions” on page 104
- 4.6, “How IBM Smarter Process for SAP creates sustained business value” on page 105
- 4.7, “IBM Smarter Process for SAP usage scenarios” on page 106
- 4.8, “Conclusion” on page 114
- 4.9, “Other IBM Software Group publications, assets, and tools” on page 116
- 4.10, “IBM Global Business Services SAP assets and tools” on page 116
- 4.11, “References” on page 116

© Copyright IBM Corp. 2014, 2015. All rights reserved.

4.1 SAP solutions as a system of engagement

Highly tailored web applications, dedicated mobile applications, and other *systems of engagement* are essential to meeting modern business expectations for market differentiation, business performance, and self-service. Conversely, enterprise resource planning (ERP) systems, such as SAP and other back-office packaged applications, provide a highly robust transactional environment.

This environment is typically optimized for transactional speed and configurability of core business functionality. Businesses today are increasingly demanding system of engagement usability and differentiation, not only in their external-facing applications, but in their core execution systems also.

Because SAP applications provide user interfaces and some concept of *process*, most SAP implementations use the SAP transactional backbone (their *systems of record*) as their default system of engagement for the vast majority of their business processes. This suboptimal approach marries an organization's business processes, and their associated business policies, to a less-than-user-friendly platform that is bound by the constraints of information technology (IT) application lifecycle management (ALM).

In most enterprises, this lifecycle typically spans three to six months, and in some cases a year or more. Yet a wide range of business changes should optimally be implemented in a matter of days or weeks, not months. Figure 4-1 shows how IBM Business Process Manager can be used as the system of engagement for the SAP system of record.

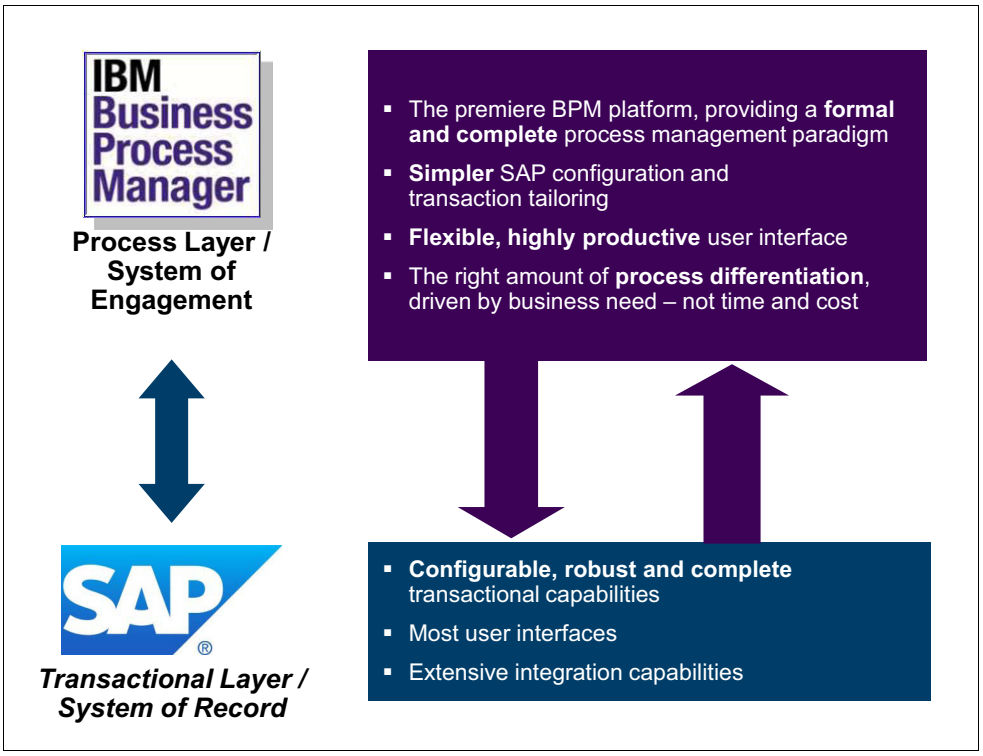

Figure 4-1 Layers of business functionality

The constraints of an IT-bound system of engagement also impose an array of technical and implementation complexities that artificially inflate the cost and risk associated with both *building* and *changing* business processes and business policy.

These complexities, and inflated cost and risk, further complicate the mission-critical task of establishing and maintaining the proper strategic alignment of business functionality provided by SAP with an ever-evolving set of dynamic business needs. Visibility, flexibility, agility, and control of SAP processes are also impaired by the sheer complexity of configuring, customizing, and maintaining business processes in what is essentially an IT-managed system of record.

Opportunities for continuous process improvement and business performance optimization are limited through the use of a system of engagement that is intrinsically bound by the IT application lifecycle. Conversely, business-led change, enabled by a flexible process layer in the system of engagement, can deliver dramatically enhanced flexibility, agility, and control over the traditional SAP implementation approach.

Externalizing at least some degree of SAP process control also has other core benefits. Many types of SAP customizations, including user interface changes, business rule changes, the addition of custom fields used primarily during the lifetime of a process instance, and transaction decomposition, can be more easily accomplished in an external process layer than in Advanced Business Application Programming (ABAP) or Java changes in SAP. The complexity of impact analysis and SAP version-related changes can be reduced also.

Embedding business rules to control business logic and process routing in an external process layer also reduces the amount and complexity of certain types of SAP configuration and customization activities, such as approval authority, skill matching, pricing, and automated credit approval. IBM Smarter Process for SAP both reduces the need for SAP customization and configuration and improves the speed with which many common types of business process and policy changes can be made.

4.2 Architecture overview

IBM Smarter Process for SAP solutions provide a comprehensive set of integration capabilities with the SAP environment, both at design time and at run time (see Figure 4-2).

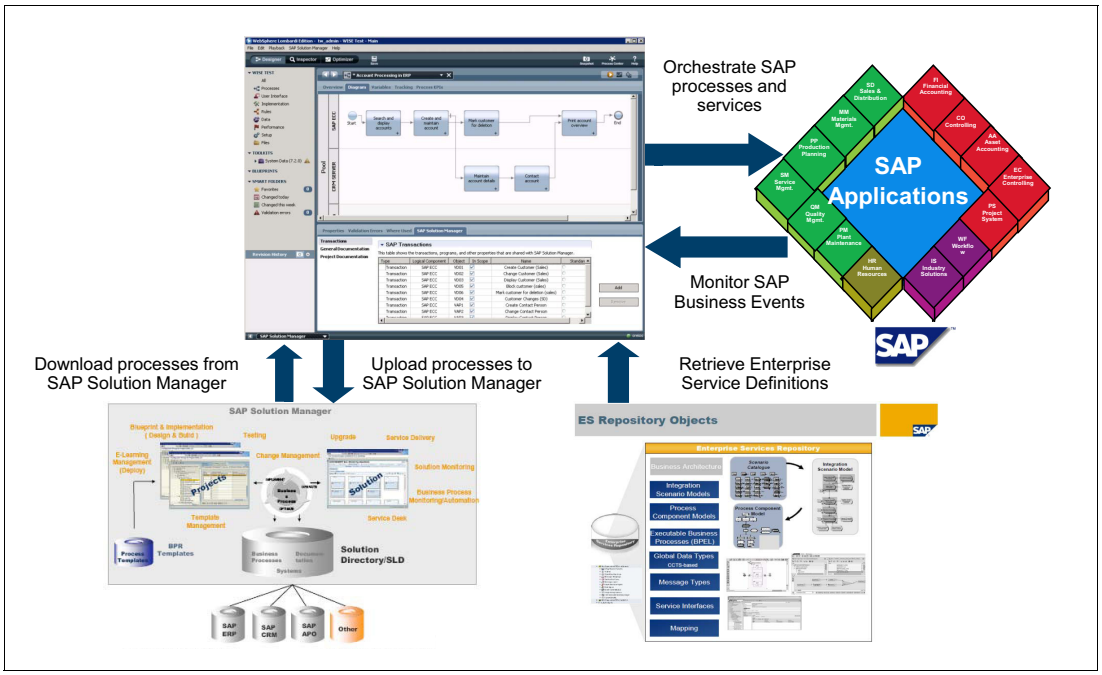

Figure 4-2 IBM Smarter Process for SAP integration capabilities

Design time integration begins with model exchange between the IBM Business Process Manager Process Designer and SAP Solution Manager. Process models can originate in the following ways:

- As business process hierarchies (BPHs) in SAP Solution Manager or in the SAP Business Process Repository (BPR)
- As business process diagrams (BPDs) in IBM Process Designer

Changes can be made in either the IBM Business Process Manager or SAP repository, and the model interchange capabilities of IBM Business Process Manager ensure highly reliable bidirectional model exchange between the two repositories.

IBM Smarter Process for SAP also provides integration with the SAP Enterprise Service Repository (ESR), which contains a library of the SAP Enterprise Services, Business Process Execution Language (BPEL) processes, and other process-related metadata useful at design time. IBM Business Process Manager can also import process models from several other modeling tools also, such as Visio and Aris.

After a model has been imported into or constructed in IBM Process Designer, the IBM Business Process Manager environment can orchestrate the following elements:

- SAP *process components*, such as SAP transactions or Web Dynpro applications
- SAP *technical services*, such as Business Application Program Interfaces (BAPIs), SAP Enterprise Services, other SAP Representational State Transfer (REST) APIs, and so on

Orchestrating SAP processes in IBM Business Process Manager uses well-encapsulated process execution steps that provide separation of responsibility between the process layer of the business and the transactional layer of the application environment.

Additionally, IBM Smarter Process for SAP can ingest and monitor SAP business events and metrics. This enables the process layer to be aware of and act upon important business activity occurring in the SAP application environment, regardless of whether an SAP process is being orchestrated by IBM Business Process Manager.

When an SAP business event is received by IBM Smarter Process for SAP, new process instances can be initiated, suspended processes can be resumed, and complex event combinations can be correlated, resolved, and acted upon.

The net effect of this comprehensive integration with SAP is to provide the complete set of design time and runtime tools required to enable business and IT teams to quickly design, manage, and deploy SAP processes without the high level of complexity typically associated with the traditional SAP process paradigm. At the same time, it provides a distinct separation of responsibilities between the transactional backbone and the process layer, to improve business agility and process flexibility.

4.3 SAP active business performance optimization architecture

The goal of implementing IBM Smarter Process for SAP solutions is to improve business outcomes. The science of analyzing, designing, and implementing these outcome improvements is typically referred to as *business optimization*. Most of the business optimization occurring in businesses today typically takes the form of *passive business optimization* (the following activities are generally optional and performed offline):

- Analyzing what needs to be improved
- Analyzing how it can potentially be improved
- Taking action to improve the business situation

Additionally, the data used to perform passive business optimization analysis is normally derived from business warehouses, such as the SAP Business Warehouse (BW) and other historical data stores. As a result, the data used for passive optimization is generally aggregate in nature, and trend-centric in scope. This leads to the discovery of scenarios that tend to point to systemic issues and generalized patterns, as opposed to solutions for specific, more granular business scenarios.

Although offline passive forms of business optimization certainly provide tremendous value to the business, a large corpus of business optimization potential remains untapped in currently running business processes. By comparison, *active business performance optimization* generally derives much of its data from currently running process instances, and tends to identify patterns that can be addressed within the time frame of a running process.

The old adage that “an ounce of prevention is worth a pound of cure” certainly applies to business practice. Active performance optimization treatments by their nature help to identify, analyze, and resolve potentially catastrophic business scenarios before they can have negative consequences. Active performance optimization can also address some of the more mundane scenarios that, although not exciting individually, often contribute in aggregate to solving some serious business problems and issues.

Several components are involved in an effective active business performance optimization architecture:

- The business needs to establish and maintain an appropriate level of actionable operational visibility.
- Business processes must be agile enough to accommodate change within the lifespan of a running process instance.
- The business must have a focus on value realization, without which the previous two elements would be without context.

Active business performance optimization requires the use of business enablers that can detect or even predict the occurrence of a negative business scenario. This capability is known as *operational visibility*. IBM Smarter Process for SAP enables the businesses to properly define, calculate, and act upon their important key performance indicators (KPIs) and performance thresholds.

It also includes capabilities that enable the business user, or the process control layer, to group instances of running processes into various subsets based on the static and dynamic parameters of the running process instance. IBM Smarter Process for SAP also provides guided optimization tools, to help with the application of both active and passive optimization techniques in the operational realm. Detection, however, is only half of the equation.

The business processes supporting the business objective must be flexible and agile enough to accommodate short-term, even one-off, solutions to a business optimization opportunity. This rapid response mechanism is known as *process agility*. This agility can be realized only if the business process, and the business rules supporting business policy, can be changed within a time frame that can capitalize upon the business opportunity during the lifetime of a running process instance.

In most traditional SAP implementations, rapid business-led change is virtually impossible. Rapid change is difficult because business processes and business rules are typically dictated by the configuration parameters of the SAP platform, and must proceed through the lengthy and costly IT application lifecycle.

Conversely, IBM Smarter Process for SAP is designed to enable rapid, business-led change that reduces the likelihood of a process, or policy changes, requiring IT lifecycle governance. Therefore, IBM Smarter Process for SAP enables short turnaround of business process changes in response to newly identified business optimization opportunities.

An important component of active business performance optimization is a focus on *value realization*. SAP implementation projects have historically focused heavily on establishing and maintaining the correct operational procedures to run the core transactional activity of the system. Accordingly, most SAP implementation projects have not concurrently introduced value realization practices as part of the implementation and, consequently, important value management discipline is deferred, sometimes indefinitely.

As indicated in Figure 4-3, IBM Smarter Process for SAP design-time and runtime integration capabilities can be included in a much broader active business optimization architecture for SAP. Business metrics, business analytics, and real-time business event sources can be combined, using the IBM Operational Decision Manager event engine as the master event source sink and correlation for SAP-specific, heterogeneous, and non-SAP sources.

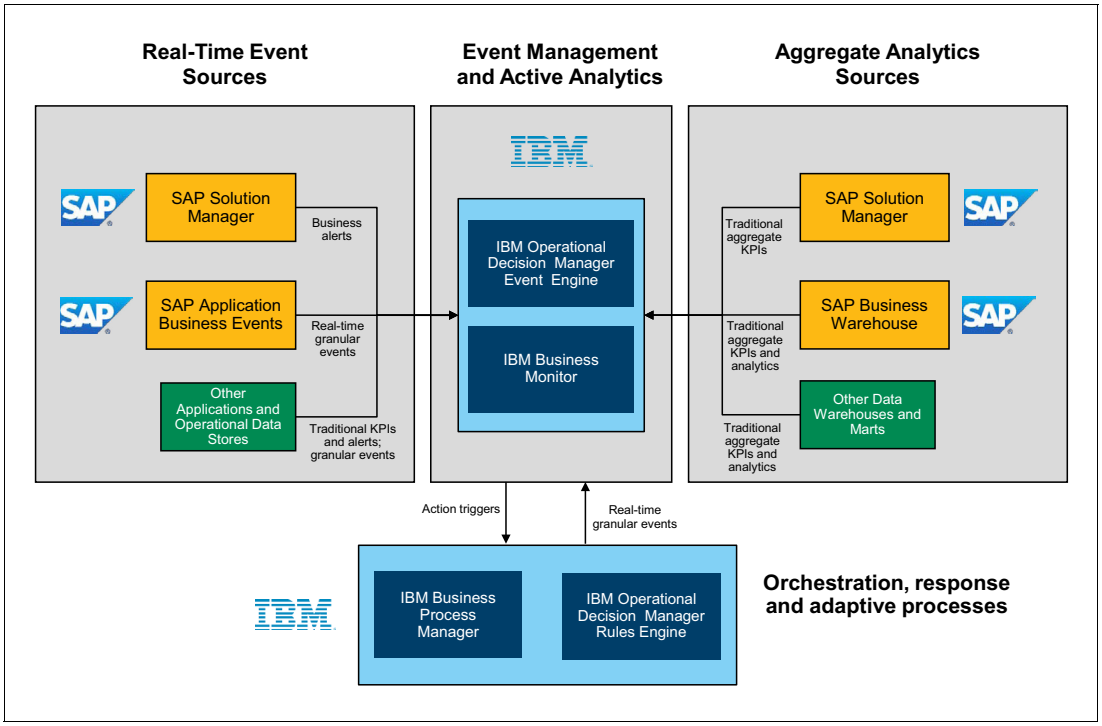

Figure 4-3 IBM active business performance optimization architecture

Using a single event management engine for all significant granular and aggregate business event sources is the only effective way to deliver an integrated view of the important events in the business, enabling you to act upon critical scenarios. Business event sources can include both SAP and non-SAP applications, data stores, and processes.

Business alerts from SAP Solution Manager, business events emitted by SAP business applications, and alerts generated by the SAP BW can easily be ingested, correlated, and acted upon. Business alerts, business events, and other relevant business information from non-SAP applications and business intelligence stores can equally be ingested, correlated, and acted upon.

Proper correlation of business events, especially when they come from disparate applications, is a complex undertaking. Most importantly, the occurrence of critical non-events must be easily detected and communicated.

The powerful inference engine found in IBM Operational Decision Manager provides tools that can be used by the business to rationalize and unify data definitions, properly define correlation schemes, and run specific correlation scenarios. Although some data and integration setup is required from IT, the bulk of the business scenario analysis, design, and maintenance can be accomplished by the business. The business requires only occasional assistance from IT to accommodate new or changing business event sources.

After an event or combination of events requiring action has occurred, the IBM Operational Decision Manager event engine triggers action in IBM Business Process Manager. Activities within the orchestrated action that has been triggered can themselves also generate events that can, in turn, be monitored, analyzed, and acted upon. The generated events become part of the original correlation scenario that triggered the initial action.

4.4 IBM Smarter Process for SAP capabilities

Whether organizations are embarking on a new SAP implementation, implementing a new SAP application module, facing an SAP version migration, or some scenario in between, IBM Smarter Process for SAP helps them to accomplish the following goals:

- Quickly know the status of key processes.
- Ensure that the process they designed is the process that is being run.
- Get real-time visibility into where workload or other bottlenecks are causing business issues.
- Effectively reroute work to less experienced workers to reduce bottlenecks.
- Know which process changes are most likely to help improve business performance.
- Quickly roll out SAP process changes.
- Quickly integrate new process workers into the business.

Documenting a business process in a modeling tool is only the first step in optimizing a business process. By using IBM Smarter Process for SAP, that same process model, used today for documentation purposes only, can be used to run, monitor, and manage a business process at any level.

Applying this external execution-centric paradigm to SAP process management facilitates an iterative solution delivery process and maintenance approach that reduces implementation time, costs, and risks. It also enables processes to be changed and enhanced at the pace demanded by the business.

Figure 4-4 shows the key interrelated process innovation capabilities for SAP application environments.

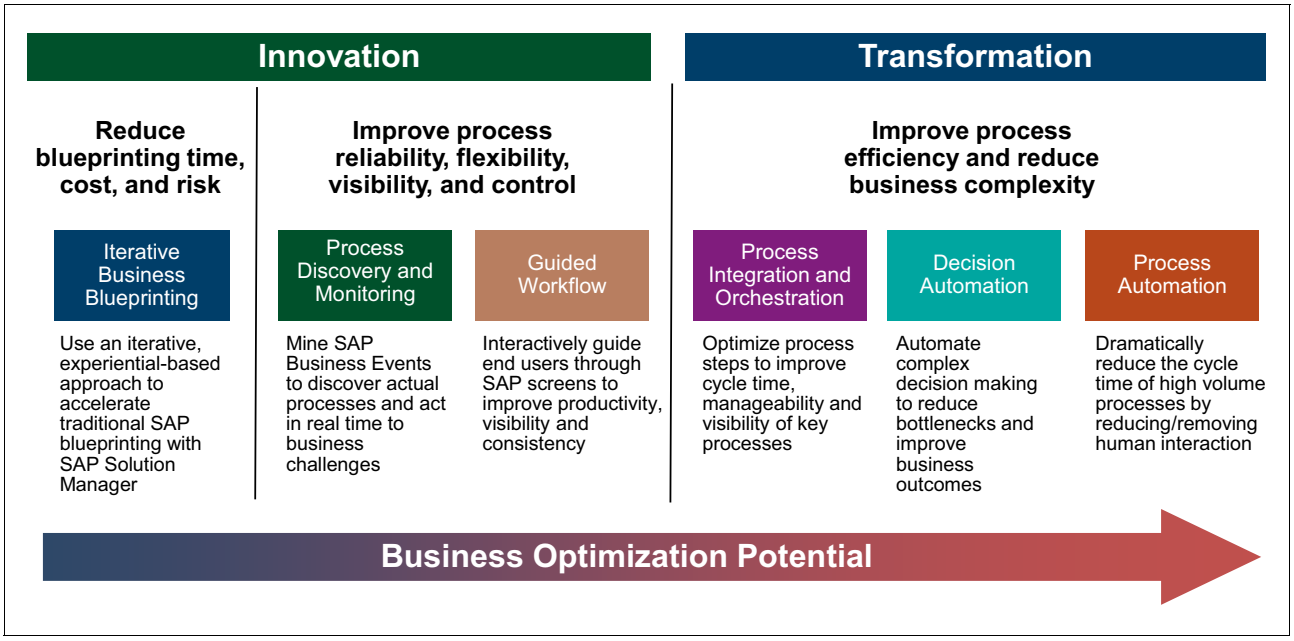

Figure 4-4 Optimizing SAP processes with IBM Business Process Manager: IBM Smarter Process for SAP capabilities

For any given process, IBM Smarter Process for SAP capabilities are normally used according to the mix required for the process type and the business optimization wanted. IBM currently provides a process affinity and value assessment workshop at no charge, to help organizations determine which of their SAP, heterogeneous, and non-SAP processes will likely benefit the most from IBM Smarter Process for SAP. The workshop also helps them determine which IBM Smarter Process capabilities are likely to help the most for each process.

To deliver these capabilities, IBM Business Process Manager provides best-in-class integration with key SAP design repositories and the SAP runtime environment. Process models can be exchanged bi-directionally and iteratively between IBM Business Process Manager and the SAP Solution Manager BPH repository.

All of the process step properties, such as SAP logical components, SAP transaction codes, documentation links, Implementation Management Guide (IMG) links and so on, can be defined directly in IBM Business Process Manager and then uploaded to SAP Solution Manager (and vice versa). Complete conflict detection and resolution capabilities are also provided, to help ensure complete process model fidelity and synchronization between IBM Business Process Manager and SAP Solution Manager.

IBM Business Process Manager enables easy orchestration of SAP transactions and Web Dynpro applications, without the need for IT development or coding. For more sophisticated process needs, IBM Business Process Manager can also enable process designers to browse, select, and automatically encapsulate and bind SAP technical services.

These service types include the SAP BAPIs, SAP Enterprise Services (web services), document flows, SAP High-Performance Analytic Appliance (HANA) APIs, and other forms of SAP technical integration. Both SAP transactions and technical services can be easily orchestrated from IBM Business Process Manager, and the resulting process flow can be used to update the BPH in SAP Solution Manager.

SAP business events, emitted whenever users or programs run SAP transactions (with or without IBM Business Process Manager orchestration), can be easily consumed, correlated and analyzed, providing a rich real-time business visibility, management, and value realization platform. When used together, these capabilities enable any organization to quickly define, change, deploy, and manage their key SAP and heterogeneous business processes.

4.4.1 SAP Solution Manager integration

A BPH that is stored in SAP Solution Manager, such as the one in Figure 4-5, can be easily imported into IBM Business Process Manager.

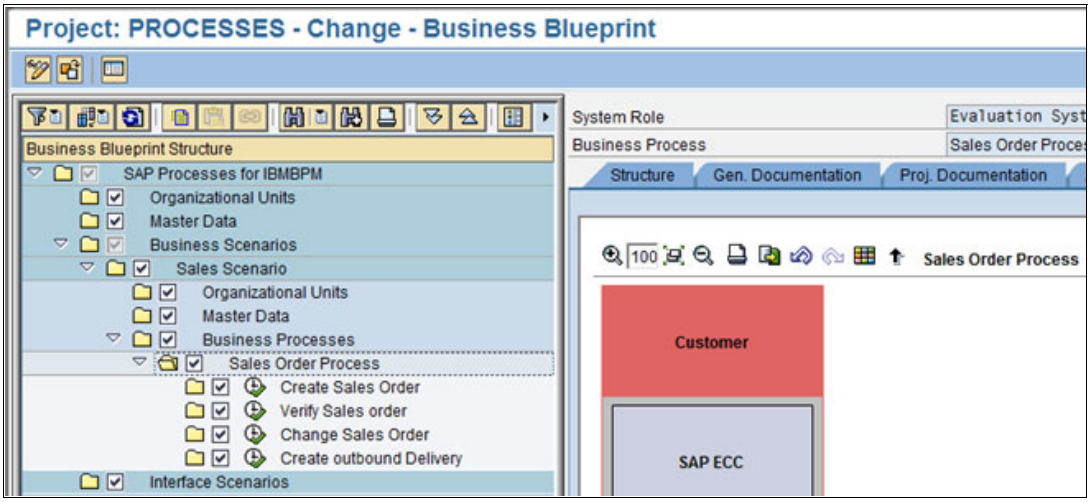

Figure 4-5 SAP Solution Manager BPH

The reverse is also true. An SAP process originally defined in IBM Business Process Manager can be exported to SAP Solution Manager as an SAP Solution Manager BPH.

Figure 4-6 shows the logon window for an SAP Solution Manager import or export interaction. Previous SAP Solution Manager connections and credentials can optionally be stored in the IBM Business Process Manager Process Center repository.

Figure 4-6 SAP Solution Manager connection

Importing an SAP BPH into IBM Business Process Manager is straightforward. Figure 4-7 illustrates the way that SAP Solution Manager Business Scenarios and business processes are selected for import. Users have the option of selecting from the following choices:

- A single business process
- Multiple business processes
- An entire business scenario
- Multiple complete business scenarios
- The entire SAP Solution Manager project

Figure 4-7 Select items to import from SAP Solution Manager

When imported, the hierarchical representation of the BPH in SAP Solution Manager is converted into a linear process flow in IBM Business Process Manager, as shown in Figure 4-8.

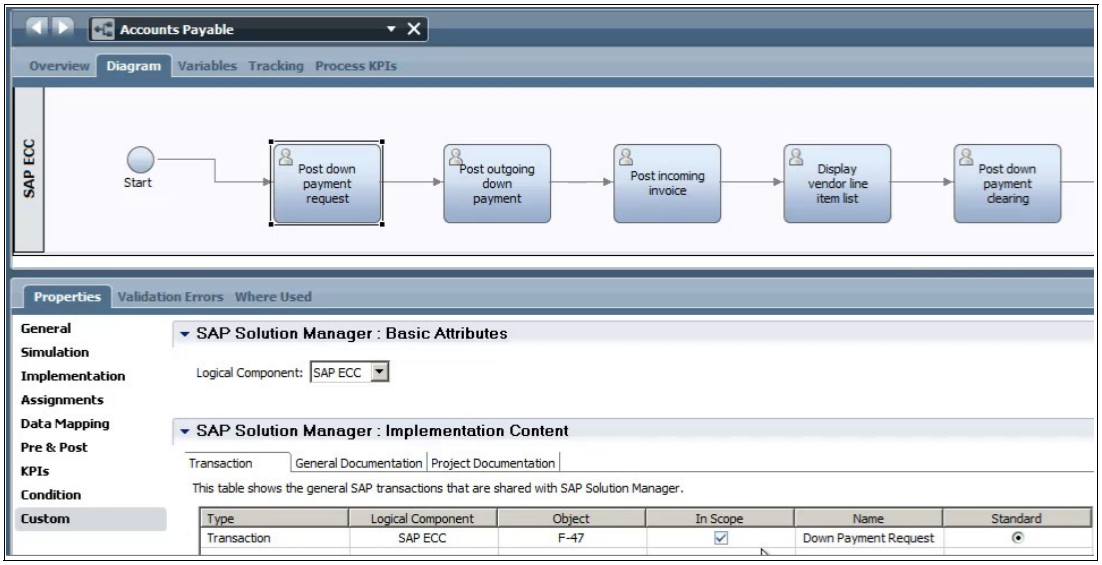

Figure 4-8 Default IBM Business Process Manager process after SAP Solution Manager import

In addition to importing BPH process steps, IBM Business Process Manager also imports the complete set of SAP implementation content that is stored in SAP Solution Manager for a given process or process step. This includes SAP transaction codes (TCODES), TCODE scoping, documentation links, and links to one or more SAP IMGs that are used for many SAP configuration activities.

In short, any BPH-related content that can be stored in SAP Solution Manager can be created, edited, and deleted in IBM Business Process Manager. This functionality provides a more convenient and complete approach for defining, refining, and communicating SAP business processes. During the import process, SAP Logical Components are mapped directly into IBM Business Process Manager “swimlanes” as a starting point for additional process refinement and definition.

For human-centric process steps, the default swimlane assignment used by the import function should be changed to reflect the actual workgroup, department, or individual process step assignments. Additional refinement of the business process definition to include manual steps, such as approval steps and escalation paths, is normally required to convert the BPH into a more complete process definition that is ready to be run.

When a business process definition has been completed in IBM Business Process Manager, or is ready for upload to SAP Solution Manager, the SAP Solution Manager export function is started from within IBM Process Designer. Similar to the import function, the user can select which processes or scenarios they want to update in SAP Solution Manager. Any conflicts or errors are identified, and the update is suspended for those processes that have errors or conflicts until these issues have been resolved.

The lifecycle management tools available in IBM Rational products can substantially enhance the governance aspects of the innovative SAP process design and execution approach available with IBM Business Process Manager. Rational tools extend IBM Business Process Manager capabilities by enabling you to map processes to requirements, test assets (such as test plans and test cases), and work items (such as plan items or user stories).

By linking process design artifacts with lifecycle management assets, the real-time planning and in-context collaboration capabilities of the Rational Collaborative Lifecycle Management (CLM) platform help project teams to apply lean and agile principles across the solution delivery lifecycle.

4.4.2 SAP Guided Workflow

Some non IBM process modeling tools (such as Aris from Software AG) also offer bidirectional BPH exchange with SAP Solution Manager. These tools are primarily modeling-only environments.

The key benefits of these tools are to help streamline the process definition phase, and to improve the quality of the documentation delivered for downstream phases of the implementation lifecycle. The end result is a small, incremental improvement to the implementation project proper, but not the kinds of transformational changes that most SAP clients would like to see in both the implementation and post-implementation phases.

IBM Business Process Manager, however, goes well beyond simply adding value to the blueprinting phase with improved documentation. With IBM Business Process Manager, every process model, regardless of its maturity level, can be run, either with what is called a *playback* or by initiating the process from the user task portal.

At a minimum, the various pathways through the process flow can be easily visualized, exercised, and analyzed. However, the most impactful benefits of IBM Smarter Process for SAP transcend simply walking through the process steps from the picture of the process flow. User interface placeholders, mock-ups, screen captures, real SAP transactions, and prototypes can easily be added for improved clarity and realism of the playback exercises.

Additionally, the standard IBM Business Process Manager process metrics, including but not limited to the following metrics, are automatically set up, calculated, and displayed by the IBM Business Process Manager environment, with no additional development required beyond defining the process:

- Total process cycle time
- Process step time
- Process step average queue size
- Process step queue size at time of execution

Arguably the most important SAP-related feature of IBM Business Process Manager is its ability to automatically generate a complete orchestration of the transactional steps needed for an SAP process.

This simple-to-use, business-led style of process orchestration, shown in Figure 4-9, requires no IT development, and is called *SAP Guided Workflow*.

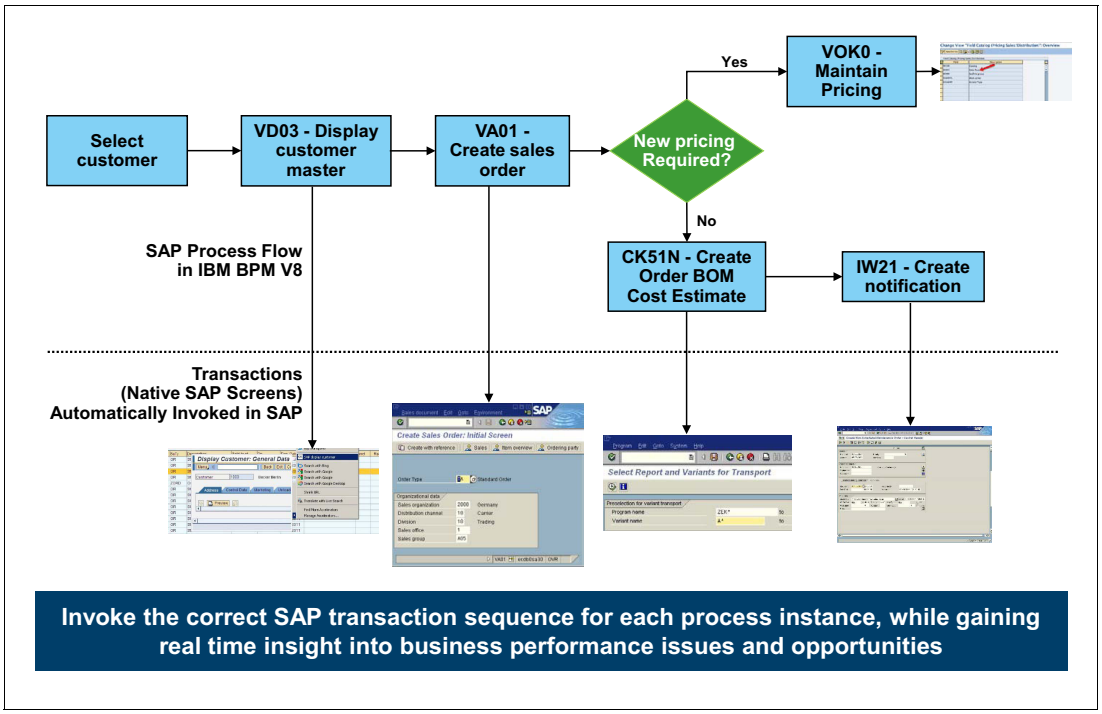

Figure 4-9 How SAP guided workflow works

As depicted in Figure 4-10, each SAP transaction in an SAP Guided Workflow process flow is exposed in the SAP Hypertext Markup Language (HTML) graphical user interface (GUI), in an iFrame in the IBM Business Process Manager user interface (UI), or in a coach view.

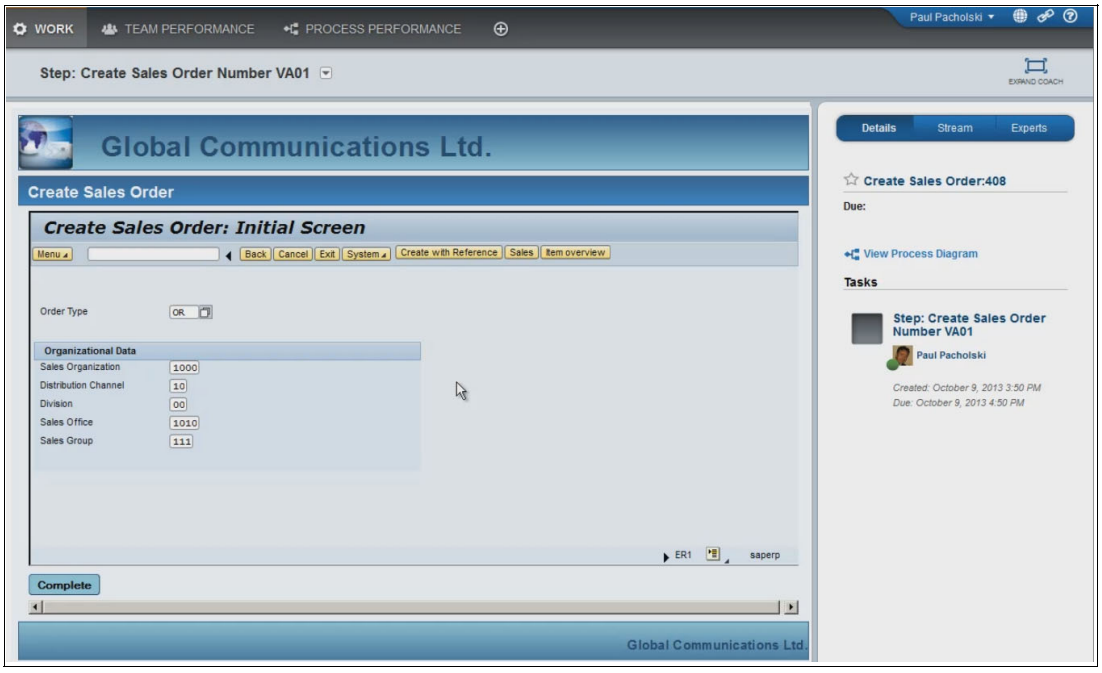

Figure 4-10 An SAP process step started with SAP Guided Workflow

IBM Business Process Manager automatically presents each SAP transaction as a standard IBM Business Process Manager process step. These capabilities are available both during blueprinting and at run time in the production environment. The full set of IBM Business Process Manager SAP Guided Workflow capabilities currently supports both SAP standard and custom transactions and SAP Web Dynpro applications.

Rather than simply providing screen mockups or static screen captures of the UI for a process step, IBM Business Process Manager starts the real SAP screens and business functionality provided by the SAP transaction or Web Dynpro application. Process designers and business participants in the process design exercise can now directly experience what the process flow will actually be, using the real SAP screens that will be used for training and in production.

This capability facilitates a more thorough process analysis and design, and enables business users who might not associate well with process pictures to fully participate in key process design and validation activities. In addition to drawing the picture of the process, IBM Business Process Manager must be able to connect to the SAP runtime system that will be delivering the SAP transactions and related business functionality. The simple parameters required to do so (System Name, Location, Client, and Port) are illustrated in Figure 4-11.

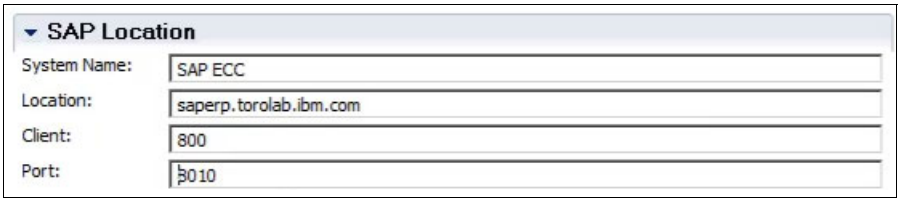

Figure 4-11 Setting the SAP runtime environment

Another key activity required to deliver a functionally complete SAP Guided Workflow process is to map values to and from SAP screen fields. This activity is required, for example, to pass an order number generated from an SAP order capture transaction into the downstream steps of the process, such as validation, pricing, picking, shipping, and invoicing.

IBM Business Process Manager enables non-IT process designers to easily retrieve data entered into any field of any SAP transaction or Web Dynpro application, and to store it in the process instance as a variable. IBM Business Process Manager also enables non-IT process designers to pass constants, process instance data, data entered into an SAP screen in a previous process step, or any other value into any field of any SAP transaction or Web Dynpro application being started as part of an SAP Guided Workflow process step.

Lastly, IBM Business Process Manager SAP Guided Workflow also enables the process designer to capture which action buttons have been activated inside of an SAP transaction. This approach helps to either confirm that the correct SAP steps have been completed, or to identify and ultimately fix the possible occurrence of unnecessary or erroneous sub-transactional activity. This bidirectional access to SAP screen data gives the process designer powerful tools to accomplish the following goals:

- Simplify work content.
- Reduce data entry workload and errors.
- Improve overall transactional accuracy.
- Reduce the need for business users to remember complex combinations of reference data to properly complete a transaction.
- Capture and rectify sources of confusion and errors.

IBM Smarter Process for SAP takes a process definition intended primarily for documentation purposes and automatically converts that definition into a fully orchestrated business process under IBM Business Process Manager control, without IT involvement. All of the important orchestration capabilities are fully enabled by the IBM Business Process Manager automated SAP Guided Workflow capability, without the need for any additional development:

- Assessing process status
- Routing work to the correct users
- Escalating problem process instances
- Starting exception or remediation processes

Figure 4-12 (with reference numbers) and Table 4-1 (with detailed explanation) illustrate how SAP Guided Workflow can be easily enhanced into a more complete process definition. That process definition can be used in all of the key process lifecycle areas: Design, prototype, test, train, and deploy into production.

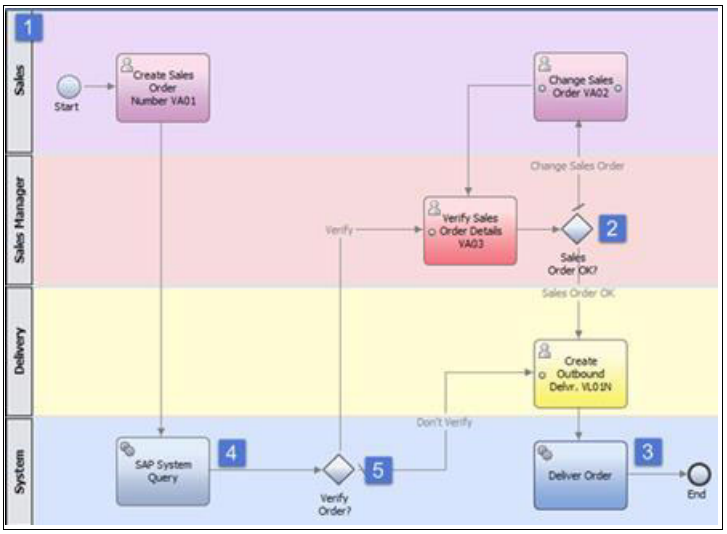

Figure 4-12 Enhancing the base SAP Guided Workflow

Table 4-1 Enhanced SAP Guided Workflow with process legend

Reference number	Description
1	Swimlanes. Each swimlane defines a team of business users that can run tasks in the given swimlane. Each team also has a manager who can supervise and manage the tasks and users using the IBM Business Process Manager Process Portal.
2	Rework loop. A decision node was added to ensure that sales managers will inspect the Sales Order (VA03) transaction. The manager can then decide if the Sales Order must be altered (VA03) or is ready for further processing.
3	This is an automated service that accepts the Sales Order Number from VA02, and the Delivery Number from VL01N, and processes them.
4	This is an automated step that accepts the Sales Order Number and returns the sales order amount for use in the decision service.
5	This is a decision node that starts an IBM Operational Decision Manager rule (using the embedded IBM Business Process Manager business rules engine). This rule determines if a Sales Order review is required (VA03), or whether to go straight to creating the delivery (VL01N).

Processes based on SAP Guided Workflow require minimal investment to build, deploy, and maintain, yet deliver the key active business performance optimization capabilities that an organization needs to get the most value from their SAP investment. For example, Figure 4-13 depicts the use of the IBM Business Process Manager happy path (best case route) analysis tool to clearly show the business why and what percentage of the time that suboptimal process instances occur.

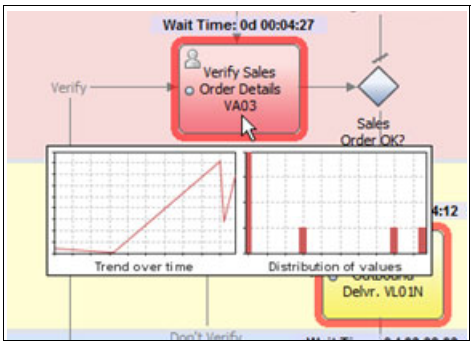

Figure 4-13 SAP process happy path analysis in IBM Business Process Manager

In this example, only 83% of the orders created actually move to downstream steps in the process. Perhaps what is most important is that fully 25% of the orders that do proceed require additional verification. By knowing these two facts, the process owner can investigate why these deviations from the happy path are occurring, and take remedial action.

One of the most common ways to continue optimizing the process is to analyze average wait times and trends to help understand the total effect of non-happy path process steps on cycle time, as shown in Figure 4-14.

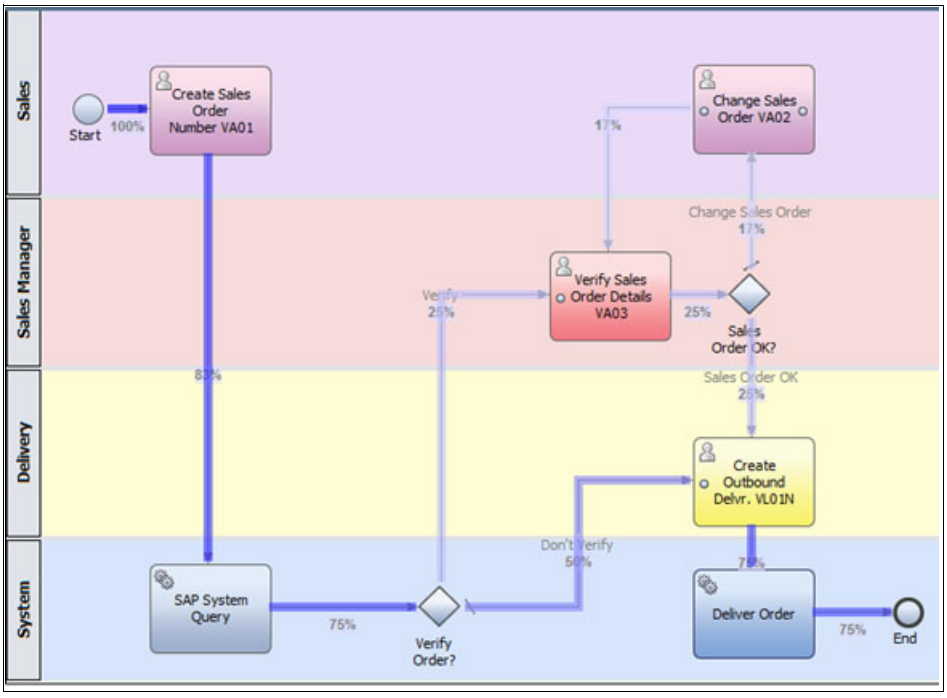

Figure 4-14 SAP process wait time analysis in IBM Business Process Manager

Because the picture of the process *is* the process, IBM Smarter Process for SAP enables the process designer to make a significant percentage of the changes identified by this type of analysis directly in the process definition itself. This is significantly faster, less expensive, and more reliable than the classic SAP process change approach.

The classic approach uses a fairly prolonged sequence of documentation, communication, training, and IT-level configuration or coding changes to effect even seemingly simple business-level changes. This is only one example of the powerful process analysis tools available in IBM Business Process Manager to help organizations improve the visibility, flexibility, agility, and control of their SAP processes.

4.4.3 Process orchestration, integration, and event management

Not every process step of an SAP or heterogeneous process can be orchestrated exclusively with IBM Business Process Manager SAP Guided Workflow capabilities. For example, some processes include simple approvals or escalation steps. These extra steps can be easily added to a process flow based on SAP Guided Workflow using traditional IBM Business Process Manager coach building and decisioning capabilities. This approach provides a fast and easy way to enhance basic SAP Guided Workflow capabilities.

For more extensive process steps, deeper technical integration with SAP and non-SAP applications is often required to consolidate and synchronize data across systems, deliver customized user interfaces, and integrate SAP process steps into broader end-to-end flows.

As shown in Figure 4-15, IBM Smarter Process for SAP offers several process integration alternatives to accommodate these more advanced SAP integration scenarios. Traditional methods of SAP technical integration, such as SAP BAPIs, SAP Enterprise Services, SAP REST APIs, and SAP Intermediate Documents (IDocs), among others, are fully supported.

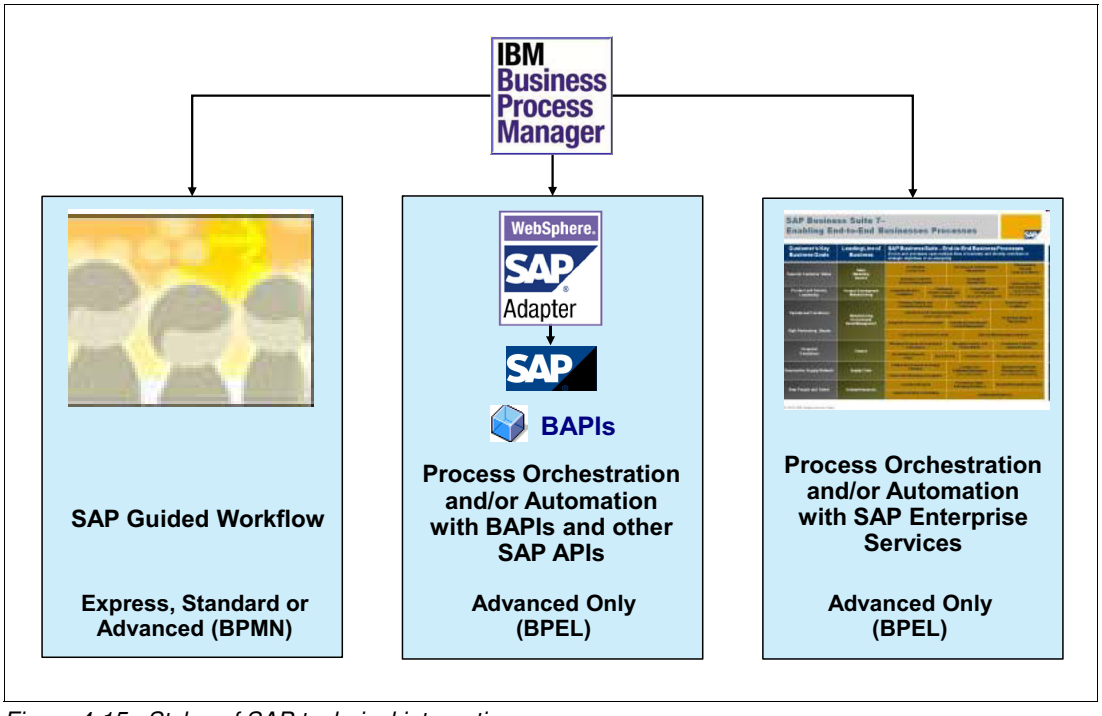

Figure 4-15 Styles of SAP technical integration

These powerful, open standards-based interfaces facilitate SAP integration by enabling the use of tools that support web services and SAP technical standards. However, this openness has historically come at the cost of complexity, because SAP technical integration generally requires coding and repetitive manual binding and encapsulation activities. IBM Business Process Manager, however, delivers advanced functionality to reduce the investment, risk, and complexity of using these powerful SAP technical integration services.

Figure 4-16 illustrates how IBM Smarter Process for SAP simplifies the invocation of SAP technical services at design time.

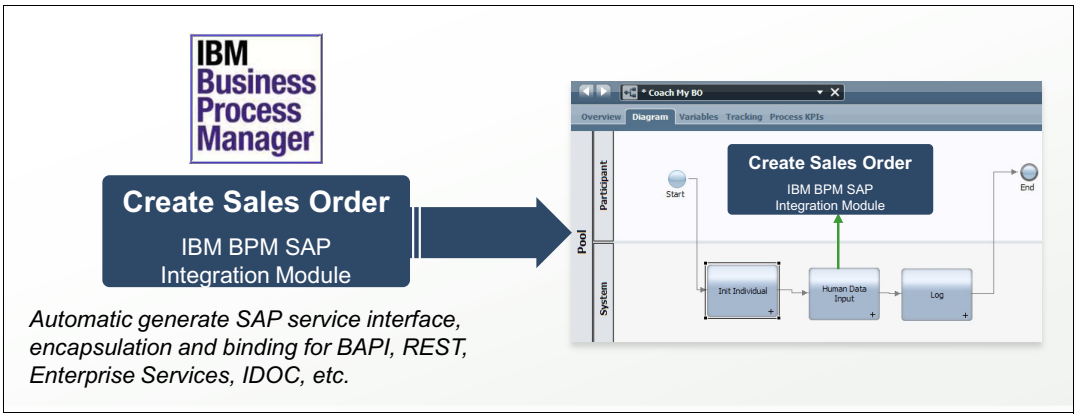

Figure 4-16 SAP technical integration in IBM Process Designer

IBM Smarter Process for SAP simplifies the technical invocation of SAP technical services so that non-programmers can select, use, and reuse these powerful services as part of a process orchestration flow, as though they were just another step in the business process. This approach enables the process designer to easily begin to optimize SAP business processes by using these services to perform the following activities:

- Automate inter-transaction activities, such as queries and lookups.
- Simplify work by reducing error-prone data entry.
- Potentially eliminate downstream activities based on the current business state of the process.

As shown in Figure 4-17, a business author or process designer defines the need for an SAP technical service interface in the context of a business process design.

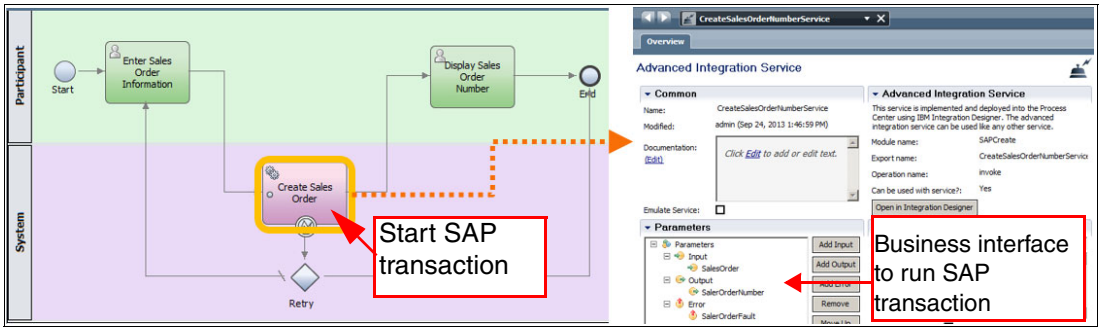

Figure 4-17 Inside the Advanced Integration Service (AIS) for SAP

The business author or process designer will typically describe the overall function needed from the service, and usually point to one or more SAP transactions that provide similar functionality. An SAP architect or technical resource then identifies which SAP technical services are required to meet the needs of the service request from the business author.

Using IBM Integration Designer and IBM Business Process Manager Advanced Edition, an IT developer then discovers the correct SAP technical services using the automated SAP service discovery tool. The developer then starts the Advanced Integration Service (AIS) implementation pattern to generate the implementation, as shown in Figure 4-18.

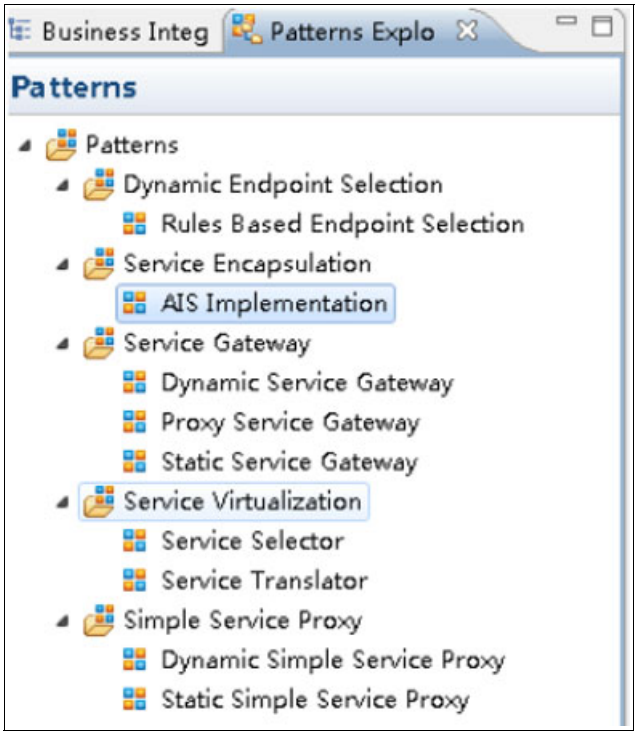

Figure 4-18 AIS pattern generator in IBM Integration Designer

Figure 4-19 illustrates how an IT developer starts the IBM Integration Designer AIS pattern to automatically bind, encapsulate, and generate the code required to run the SAP technical service as a simple process step in IBM Process Designer.

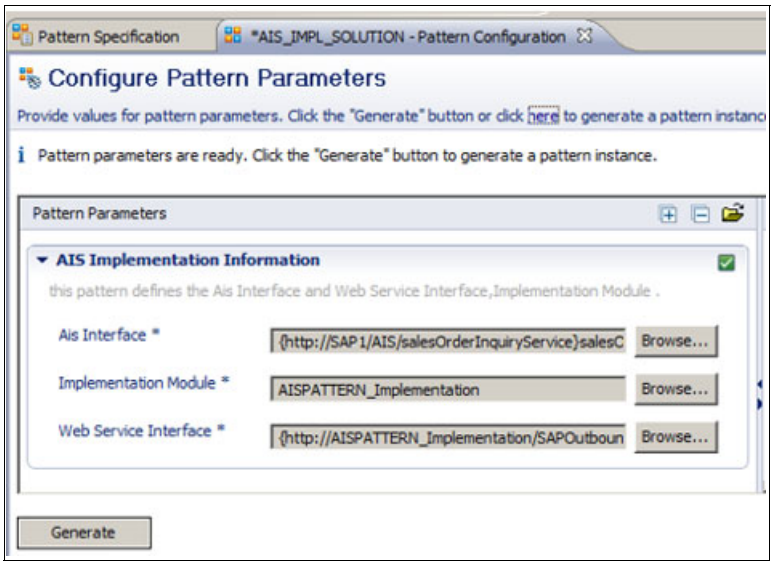

Figure 4-19 Selecting parameters for automated AIS generation for SAP

The generated mediation pattern includes service invocation, fault handling (both business and technical), and data mapping (input data, output data, and faults), as shown in Figure 4-20.

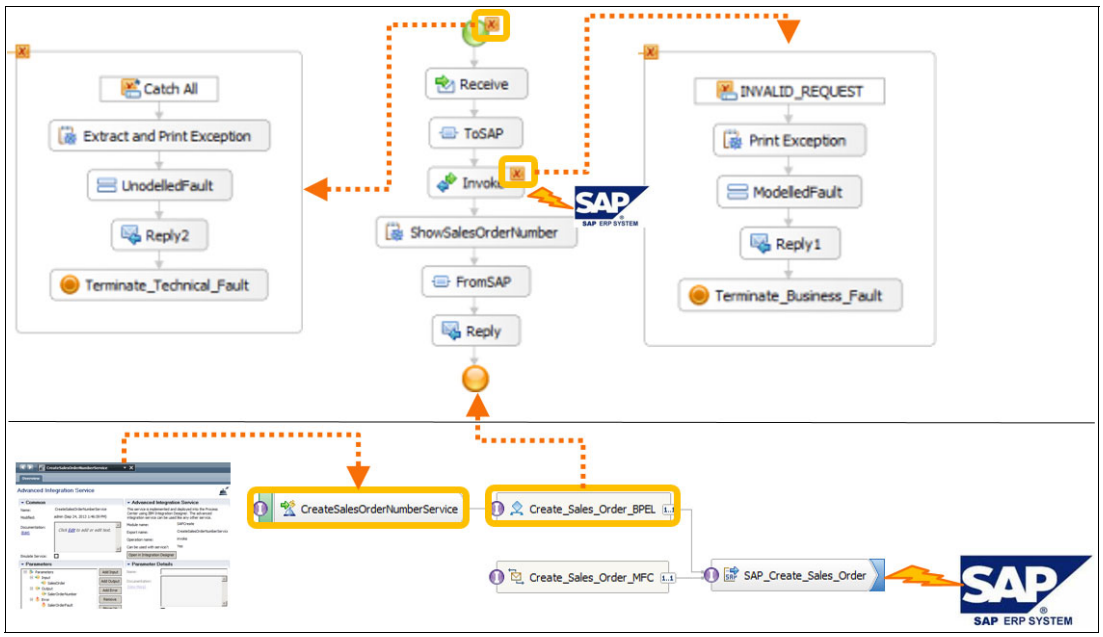

Figure 4-20 Default error handling in the AIS for SAP

The final step is for the IT developer to use the standard IBM Integration Designer graphical data mapping tools to complete the data mapping between the AIS and the SAP technical service.

As Figure 4-21 shows, the IT developer can also trim any unwanted output from the SAP technical service to streamline usage of the AIS for typical business scenarios composed in IBM Process Designer.

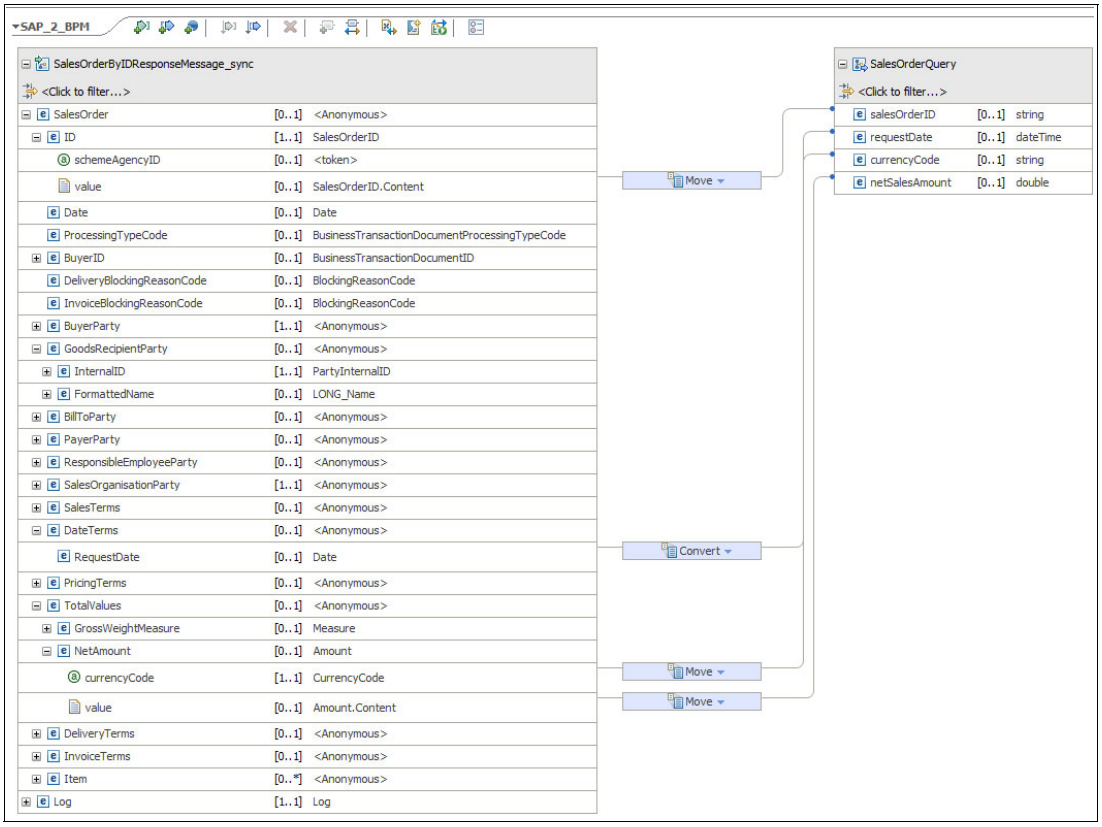

Figure 4-21 Reducing AIS for SAP interface complexity

This approach masks the complexity of SAP technical integration for the business process designer. It also enables process designers to seamlessly integrate both human-centric and technical interfaces, all in the same process model and in the same manner.

Compared with manual approaches to using SAP technical integration, IBM Smarter Process for SAP saves substantial time and money, reduces project risk, and encourages the adoption of active business performance optimization. IBM Business Process Manager SAP technical integration pattern technology also automatically applies a consistent set of leading practices for SAP technical integration, further reducing the investment required to build, use, and maintain the use of powerful SAP technical integration capabilities.

4.4.4 Process discovery and monitoring

Because of the inherent limitations of the documentation-centric approach typically used to implement SAP solutions, SAP business process documentation is quickly outdated and out of sync with the actual processes being used inside the business. Fortunately, SAP provides its customers with a set of business events that can be analyzed offline to better understand the actual processes in use and their performance characteristics.

The data treatment techniques and discovery process required to gain this insight are quite complex, so most organizations do not use SAP Business Events for this purpose.

Figure 4-22 shows how IBM Smarter Process for SAP offers a distinctive set of capabilities that help business users understand their actual SAP processes by showing the SAP transaction flows in a process view, without requiring orchestration of the SAP processes.

In the bottom part of the screen, users can see a picture of their SAP process, with key process and business content data and analytics available for each process step, and for the process at large. Standard process analytics include but are not limited to the following items:

- Total cycle time
- Activity time
- Activity queue time
- Activity queue size

Analytics on business content can include anything that is relevant to the process step, and is available from either the corresponding SAP business object or from the amalgamation of business data that has been modeled and stored in the monitor model.

The top half of the screen provides a tabular view of the SAP transaction instances that have passed through the process step selected on the bottom half of the screen. This tabular view can have multiple levels of detail, such as order header, order detail, shipment header, and shipment detail. It can include both business data and process analytics. The business user can then drill down and, across the dimensions of the table, filter results and take action based on the data presented.

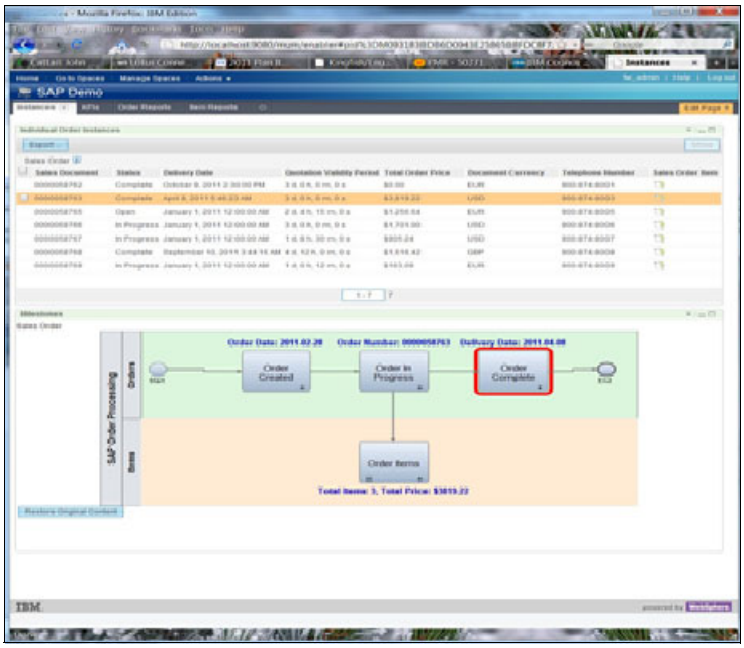

Figure 4-22 SAP process characterization in IBM Business Monitor

IT developers build SAP Business Event listeners (authored using the inbound capability of WebSphere Adapter for SAP Software) that listen for a specific SAP Business Event. The listener then retrieves relevant SAP Business Object information, and emits Common Event Infrastructure (CEI) Common Base Events that then deliver a complete packet of business event information to IBM Business Monitor for analysis and action.

As shown in Figure 4-23, IBM Business Monitor then receives and correlates these events.

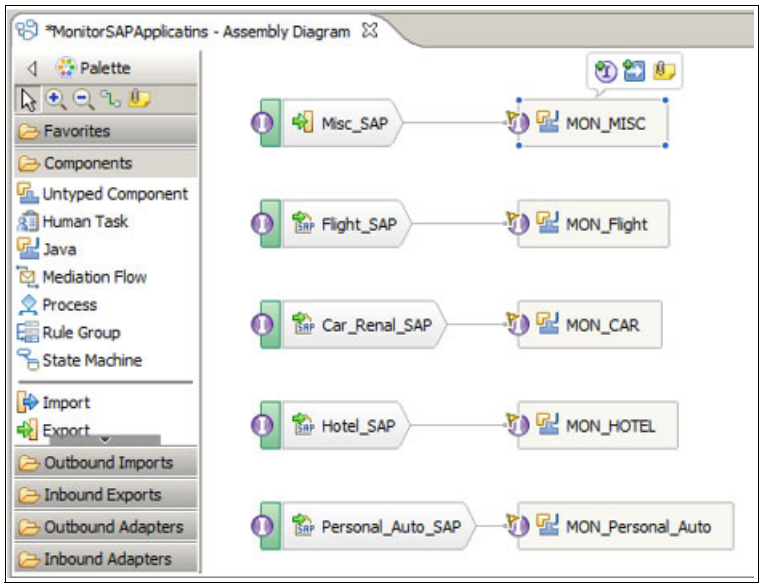

Figure 4-23 Configuring SAP Business Event listeners in IBM Business Monitor

IBM Business Monitor displays the events in web-based dashboards, such as the automatically generated milestone diagram shown in Figure 4-24. This diagram can be automatically generated by IBM Business Monitor based on a global view of the monitor model, and can quickly depict SAP transaction flows without the amount of development required to deliver the content and outputs listed in Figure 4-22 on page 96.

The milestone map approach is appropriate where a quick view of the process and its process analytics is required without the need for the more advanced capabilities, such as the tabular view, action management links, and so on.

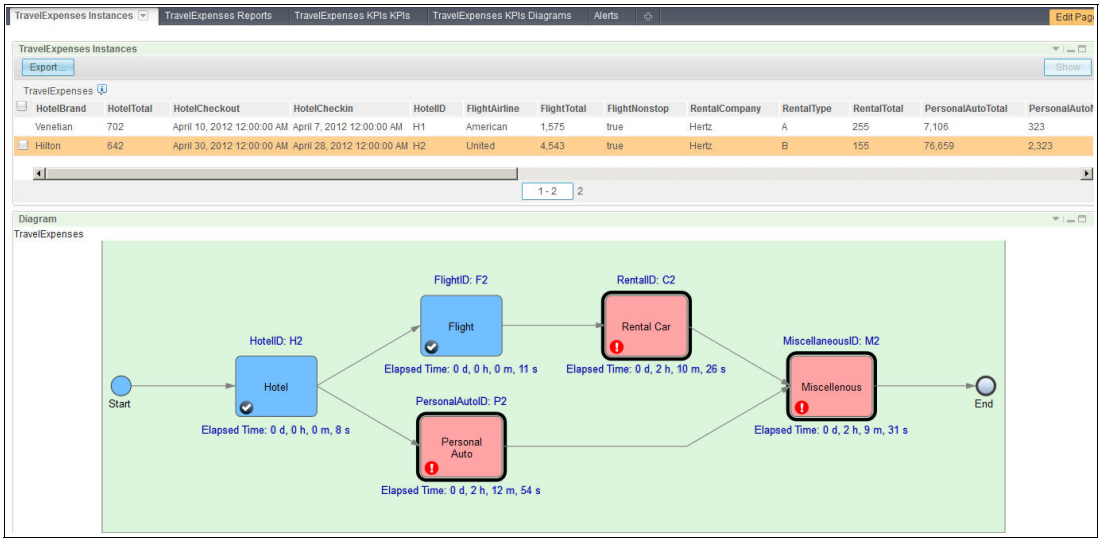

Figure 4-24 Automated SAP milestone mapping in IBM Business Monitor

Even before you consider orchestrating an SAP process in IBM Business Process Manager, this powerful capability enables you to empirically characterize key aspects of your SAP and heterogeneous processes:

- Total cycle time
- Process step activity time
- Lag time between process steps
- Queue metrics across the process

With this information, you can easily gain the empirical insight required to identify key areas for process improvement. This functionality also facilitates a qualitative understanding of the business flows, including sequencing, out of sequence occurrences, exception paths, exception metrics, and so on.

These same capabilities in IBM Business Monitor that deliver passive process analytics and characterization, also provide easily configurable real-time active monitoring for SAP. After the process scenario is properly defined in the tooling, IBM Business Monitor automatically configures and manages the data stores, triggering mechanisms, and so on, that are required to define, visualize, operationalize, and institutionalize the performance management layer built on KPIs and service level agreements (SLAs).

Important business measures, and their performance thresholds, can then be used to automatically trigger preventive and reactive remedial processes. Therefore, many of the active optimization capabilities of IBM Smarter Process can be imparted to virtually any SAP process, regardless of whether it is being orchestrated by IBM Business Process Manager.

All of these capabilities are fundamentally non-intrusive to, and decoupled from, the SAP environment, enabling easy setup and low impact on the SAP platform itself.

IBM Business Monitor enables organizations to perform the following tasks:

- Monitor metrics, business situations, and events in real-time.
- Deploy customizable dashboards to ensure targeted, relevant information.
- Feed and correlate alerts with business event processing for enhanced pattern visibility.
- Interact directly with processes in real time.
- Predict future values of KPIs based on historic and cyclic trends.
- Trigger alerts when predicted values indicate a problem detection.

Figure 4-25 shows an example of a monitoring dashboard and illustrates the key capabilities provided.

Figure 4-25 IBM Business Monitor dashboards for passive optimization

The reference numbers in Figure 4-25 highlight key capabilities of the monitoring dashboard and Table 4-2 provides the corresponding description.

Table 4-2 Monitoring dashboard key capabilities shown in Figure 4-25

Reference number	Description
1	Identify trends, forecast events, make smart choices.
2	Understand up-to-the-minute business performance by monitoring KPIs.
3	Detect and respond rapidly to change.
4	Continuously improve key business processes.
5	Rebalance human workload dynamically.
6	Customize dashboards easily.
7	Use mobile devices.

4.4.5 Iterative business blueprinting

Most SAP customers and consultants continue to use the same traditional waterfall approach for SAP business blueprinting that has been the hallmark of SAP implementations for at least the past 20 years.

As depicted in Figure 4-26, general-purpose documentation tools, such as Microsoft PowerPoint or Visio, or a modeling tool, such as Aris for SAP process blueprinting, are used to define high-level SAP process flows (generally in conjunction with the SAP Solution Manager BPH and SAP BPR).

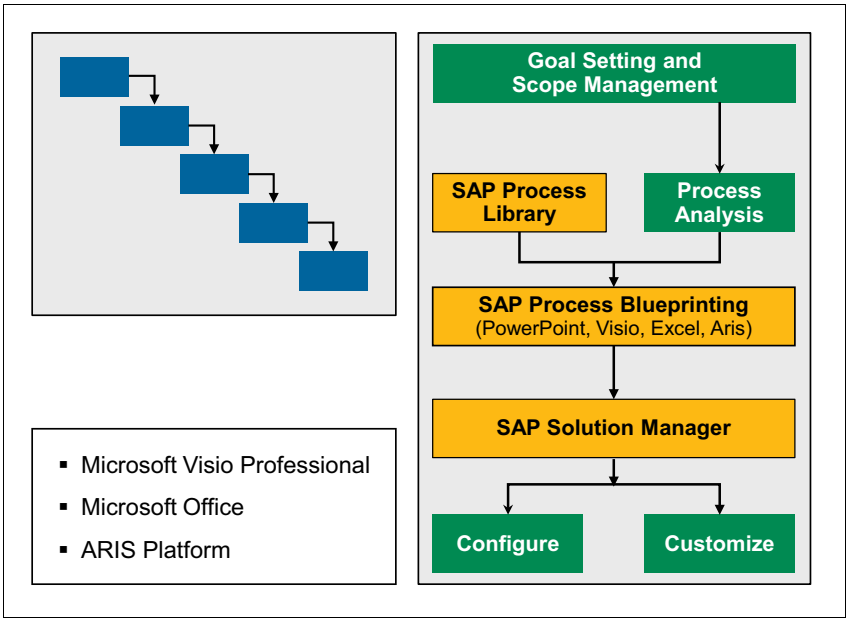

Figure 4-26 The traditional waterfall SAP business blueprinting approach

Business domain and process area experts are then engaged at various points throughout the six or so months that span most SAP blueprinting cycles, primarily by analyzing the process documents produced by teams using these tools.

With this approach, business stakeholders are forced to approach their SAP business blueprinting role analytically, which limits the range of business experts and process design techniques that can be used to help shape the new SAP business processes. Although sequentially linking a cascading set of well-defined SAP business blueprinting activities is logical, the deficiencies of the waterfall *analyze first, then design, then build* approach become apparent when alternatives are considered.

IBM Business Process Manager uses the iterative, playback-based process design illustrated in Figure 4-27 to help make the definition and improvement of SAP and heterogeneous processes transparent, while accelerating the blueprinting process.

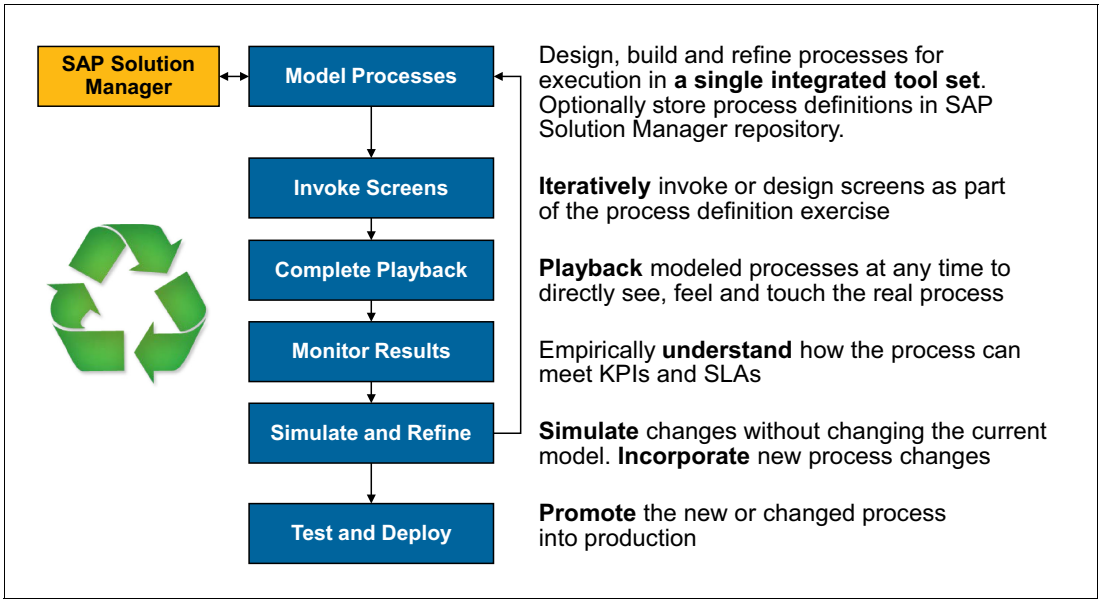

Figure 4-27 The iterative SAP blueprinting approach using IBM Smarter Process

As opposed to the traditional waterfall documentation-centric approach, iterative playback-based process blueprinting integrates *live* SAP user interfaces and process execution, using SAP Guided Workflow, into the blueprinting process. This enables process workers and other business experts to directly develop and test an SAP process hands-on.

The blueprinting process becomes experiential, rather than analytical. Feedback is immediate, and most changes to the SAP process can be made directly in the workshop itself. The process flow can be tested then and there, reducing the time span of blueprinting from months to days.

In turn, this approach dramatically reduces blueprinting costs and risks. Using IBM Business Process Manager, the processes as defined in the SAP business blueprint can also be directly run in the test, pre-production, and production environments. This ability to directly run the process reduces SAP customization while delivering the improved visibility, management, and control afforded by process orchestration.

Playbacks enable each stakeholder to participate *directly* in the blueprinting process. The ability to see, feel, and touch the process as it is running provides the following benefits:

- Delivers a richer, superior design experience
- Enables a broader range of participants
- Enables stakeholders to participate as their schedule permits
- Encourages a progressive realization build approach
- Decreases the time required to blueprint a process

4.4.6 Decision automation

Many important business decisions, such as how to allocate constrained inventory or how much additional credit to extend to a repeat customer, are often made inline, while the process is in midstream. Today, many of these decisions are driven by policy documents, email guidance from upstream management, offline tasks, or casual collaboration. This leads not only to inconsistent application of business policy, but also to significant delays in changing these important decision-making parameters as business needs require.

IBM Smarter Process for SAP offers a powerful set of integrated tools to help organizations automate and enforce these important business policy parameters, either inline with the process itself or offline. Business policy changes can be made at the speed required by the business to optimize business performance, respond to dynamic business conditions, and improve the consistency of routine decision making.

IBM Smarter Process for SAP decision management capabilities help to optimize business processes inline:

- Automatically routing process steps to the correct workgroup or individual
- Automatically eliminating optional steps where business policy and business conditions permit
- Enabling a single process definition to fulfill multiple process variants by automatically selecting the correct process steps that form the variant at run time
- Making processes dynamically adaptable to changing business conditions by automatically selecting and routing process steps at run time, depending upon any set of process instance and business environment attributes
- Automatically starting companion processes or remedial action based on the correlation of complex business events

Process designers have the choice of using the integrated business rules capabilities that are part of IBM Business Process Manager or can optionally deploy business rules as services using the more advanced capabilities of IBM Operational Decision Manager.

The advanced capabilities in IBM Operational Decision Manager enable business rules to be encapsulated and reused across multiple business processes. It also provides a more comprehensive set of rule authoring paradigms such as decision tables, decision trees, and Microsoft Excel spreadsheets.

Business rules developed in IBM Business Process Manager are compatible with IBM Operational Decision Manager, enabling process designers to easily migrate business rules from IBM Business Process Manager to IBM Operational Decision Manager whenever these more advanced capabilities are required.

Figure 4-28 shows an example of business rules encapsulated as decision services that are, in turn, embedded as an integral part of the business process.

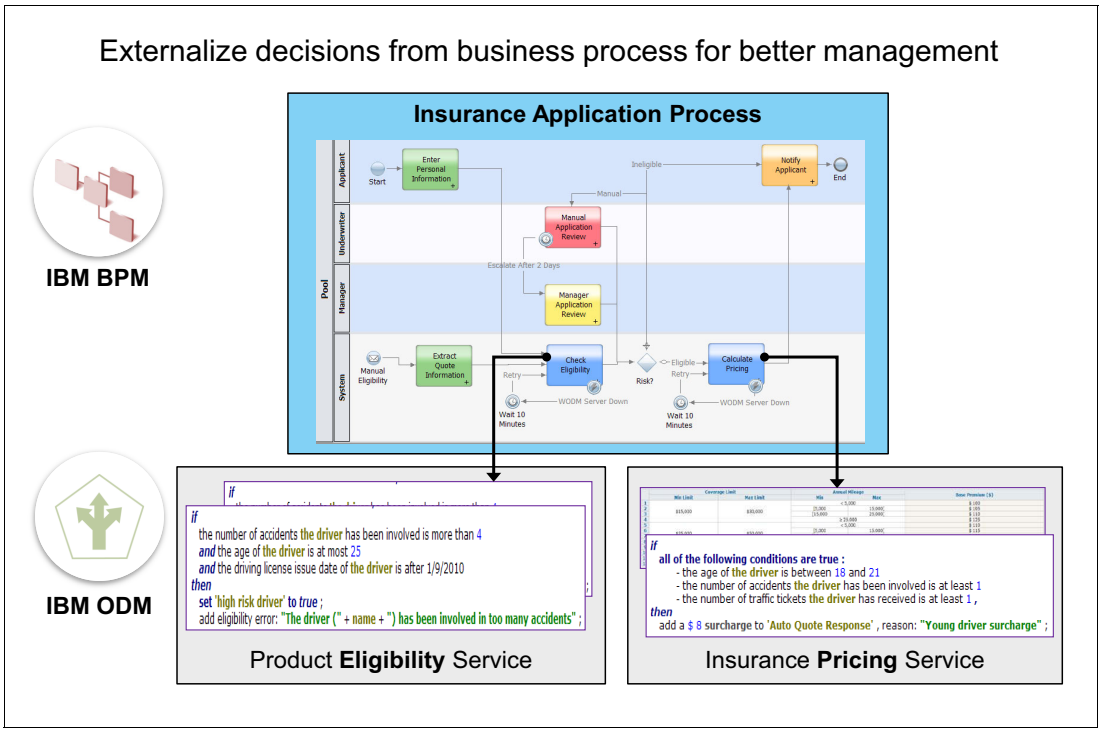

Figure 4-28 Encapsulation of decision logic using decision services

4.4.7 Process automation

IBM Smarter Process for SAP provides the full set of capabilities required to automate any aspect of the business processes in whole or part. Powerful technical integration for SAP and other business applications, data assets, and third-party services reduces the time and cost required to create and maintain process integration and automation capabilities. The same technical integration capabilities can be used to automate a single process step.

These capabilities can also be used to create detailed and powerful end-to-end process flows that span multiple applications, departments, and organizations, such as the example shown in Figure 4-29.

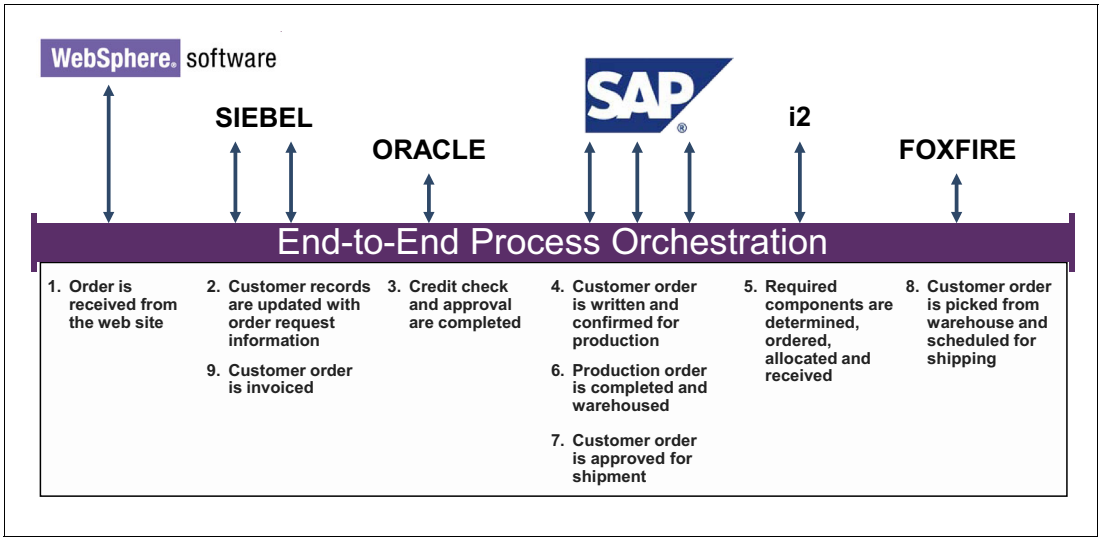

Figure 4-29 IBM Smarter Process end-to-end process orchestration

Figure 4-29 shows how IBM Business Process Manager can integrate with the full set of business applications and technology platforms that make up this example Order to Cash scenario to deliver a cohesive, well-managed process for the business.

For some of the platforms in this example, such as Siebel, Oracle, and SAP, IBM provides application integration adapters to deliver the technical integration necessary for end-to-end process integration and orchestration.

For other platforms, such as Foxfire, an adapter development toolkit is available from IBM to help you quickly develop the technical integration capabilities that you need. For virtually every significant end-to-end process in the business, IBM Business Process Manager helps to coordinate tasks across platforms, improve collaboration inside and outside the enterprise, reduce cycle time, improve productivity, and enhance business outcomes.

4.5 IBM Smarter Process for SAP products and solutions

This section lists several key IBM products that can be used when implementing IBM Smarter Process for SAP scenarios:

- IBM Business Process Manager
  - Orchestrates, manages and optimizes virtually any business process (SAP, non-SAP, and heterogeneous).
  - Enables bidirectional exchange of the SAP BPH and implementation content with SAP Solution Manager.
  - Provides SAP Guided Workflow, a powerful, business-led set of capabilities to orchestrate SAP transactions and Web Dynpro applications without IT development. Includes bidirectional integration with SAP screens.

  - Delivers easy-to-use SAP technical integration to help use SAP BAPIs, Enterprise Services, IDoc flows, REST APIs, and so on, without the extensive coding required with traditional approaches. This integration can reduce or eliminate manual steps between and inside SAP transactions, and help to integrate and automate end-to-end process flows.
  - Provides a powerful set of process-level work and capacity management tools, along with a rich set of process analytics and an innovative guided optimizer for active business performance optimization.
- IBM Operational Decision Manager
  - Automates and assists critical business decisions inline by delivering industry-leading business rules and event correlation technology.
  - Packages business design logic as decision services for reusability and consistency across the enterprise.
  - Delivers automated task routing and prioritization.
  - Provides automated process step selection in multi-variant and dynamic business processes.
  - Facilitates the selection of SAP runtime environments in multi-instance SAP processes.
  - Helps to quickly modify business policy for agile SAP processes, while reducing the amount and complexity of SAP configuration.
- IBM Business Monitor
  - Provides real-time operational visibility of any SAP, non-SAP, or heterogeneous process, even without Business Process Manager orchestration.
  - Delivers SAP process characterization for passive and active business insight.
  - Automatically calculates KPIs, SLAs, and other business measures.
  - Automatically starts orchestrated or passive responses to business measure threshold violations or negative trending, for active performance optimization.

4.6 How IBM Smarter Process for SAP creates sustained business value

IBM Smarter Process for SAP creates business value throughout the SAP lifecycle. For net new *green field* SAP or SAP module implementations to SAP version migrations and everything in between, IBM Smarter Process for SAP can help deliver the following benefits:

- Increased business velocity
- Business optimization at the pace required by the business
- Improved alignment of processes to business strategy
- Improved managed value realization
- Closed-loop continuous process improvement
- Reduced SAP customization
- A comprehensive platform for organizational transformation
- Reduction in manual work, redundant tasks, and data errors
- Faster, more consistent issue resolution
- Improved control over business management parameters
- Reduced blueprinting time, cost, and risk
- Improved process reliability, flexibility, visibility, and control
- Improved process efficiency and reduce business complexity

4.7 IBM Smarter Process for SAP usage scenarios

IBM Smarter Process for SAP capabilities are typically combined in different ways, depending on the usage scenario. These capabilities can be applied across the entire SAP project lifecycle from concept, through blueprinting and realization, to post implementation business optimization scenarios. This section shows how IBM Smarter Process for SAP can add value in several different project and business scenarios, both inside an SAP implementation project and after go-live.

Figure 4-30 shows the two outer boundaries of an SAP implementation timeline: the *new* SAP implementation and an SAP *version migration*. Both of these major SAP lifetime events are ideal places to consider the application of IBM Smarter Process for SAP capabilities, because business processes, business policy, and technical and application infrastructure are explored and changed extensively in these project types.

Often, IBM Smarter Process for SAP can be inserted into a new SAP implementation (including instance consolidation, harmonization, global template rollout projects, and so on) with no negative effect on either the cost or timeline of the project as compared to the same project scope without IBM Smarter Process for SAP. SAP version migrations that include a functional upgrade can likewise be excellent candidates for the introduction of IBM Smarter Process for SAP capabilities.

Between these two points, IBM Smarter Process for SAP can be applied more selectively to any number of end-to-end process scenarios, or used to innovate and optimize some of the more granular levels of a business process. It is also an ideal platform to extend or modify SAP business functionality, because of its loose coupling (but extensive integration) with the SAP environment. Stable SAP implementations are also ripe for either the initiation or expansion of an IBM Business Process Manager or SAP Center of Excellence (CoE).

These can serve as the mechanism to find, incubate, and deliver business optimization based on IBM Smarter Process for SAP.

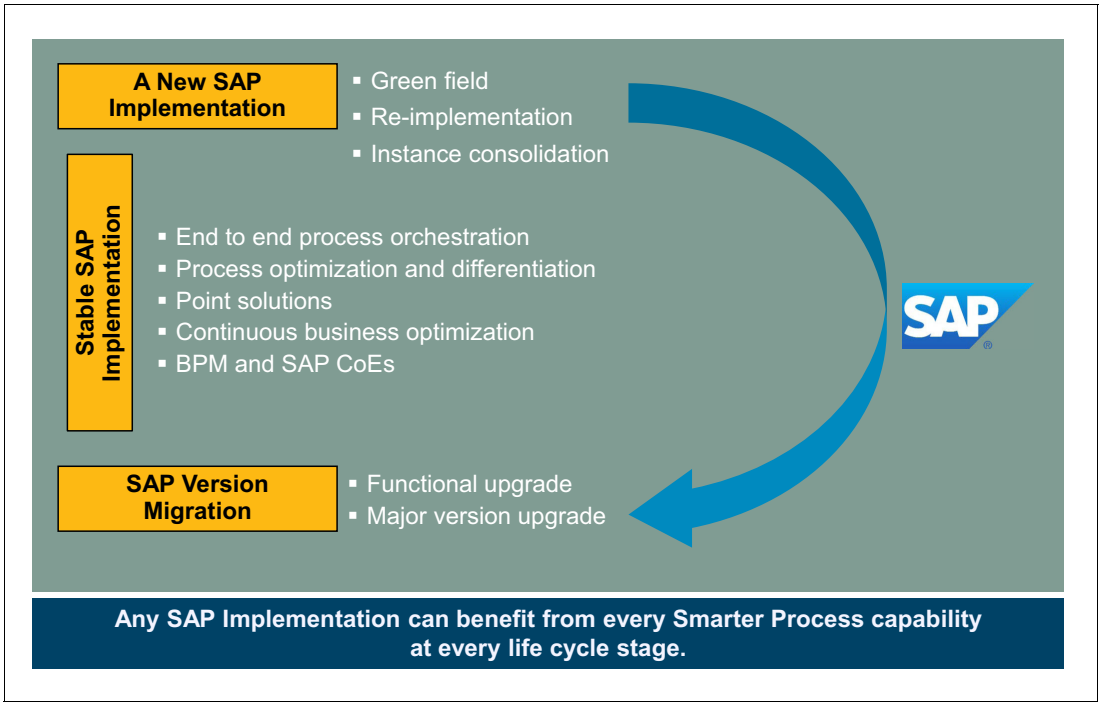

Figure 4-30 IBM Smarter Process for SAP project types

4.7.1 IBM Smarter Process for SAP in the phases of an SAP project

The three major phases in the life of an SAP implementation project are as follows:

- Initial implementation
- Stable implementation
- Functional upgrade or re-implementation

IBM Smarter Process for SAP can add substantial value to each of these phases. Although the manner in which organizations can use IBM Smarter Process for SAP varies somewhat from phase to phase, the capabilities, design philosophy, and core value proposition are more or less the same. This section explores how IBM Smarter Process for SAP can be used throughout the SAP project lifecycle.

Iterative business blueprinting

The start of an SAP implementation is an ideal time to begin using IBM Smarter Process for SAP capabilities. Business processes will be designed entirely differently knowing that the process orchestration, business activity monitoring, and value realization capabilities available in IBM Smarter Process for SAP will be available for the SAP deployment.

As described in 4.4.5, “Iterative business blueprinting” on page 100, SAP business blueprinting has been performed fundamentally the same way for the last 20 years or more. This classic implementation approach is analytical in nature, uses IT engineering methods, and is deployed using a waterfall methodology. These kinds of activities are foreign for the average business user who must *cross the chasm* between pictures of business activities and those business activities themselves.

SAP Guided Workflow offers new possibilities for business blueprinting. Because of the immediate playback capabilities of SAP Guided Workflow, business users participating in blueprinting workshops can be guided not only through a description of the steps of a process, but can actually see, feel, and touch the SAP screens that will be used.

Rather than just hearing about the kinds of information that need to be entered into an SAP screen, blueprinting workshop participants can see firsthand the nature and content of work required for a particular process step, in addition to the overall process flow and composite workload.

This level of understanding has a profound effect on many aspects of detailed process design, and helps provide better alignment between what the business needs and what SAP can provide. This approach removes the historical barrier that has existed between engineering-centric blueprinting artifacts and the business user. It enables the line of business to participate meaningfully in the business blueprinting process, and the progressive realization of their business solution.

IBM Smarter Process for SAP reduces the time, risk, and cost of SAP business blueprinting by reducing the size and complexity of the artifacts required to properly design and implement an SAP solution.

IBM Business Process Manager orchestration definitions are by their nature rich with process metadata. They contain a high percentage of the metadata that is normally defined in SAP blueprinting documentation, such as the Process Definition Document (PDD), the process model, and the functional definition document.

Consolidating all of this metadata directly into the same process model that is being used for iterative workshops results in a reduction of verbiage needed to describe key aspects of the process and improved transparency of the process definition.

Blueprinting productivity improvements ranging from 10% to 25% are achievable using SAP Guided Workflow in conjunction with iterative, experiential process design techniques. Most importantly, iterative blueprinting begins the journey of business-led change, and enables the business to directly contribute to, and manage many of, the important aspects of properly defining and optimizing a new or revamped SAP business process.

Progressive solution realization

The program benefits achieved by using IBM Smarter Process for SAP for business blueprinting continue in the progressive realization phase. The iterative process that is used to define an SAP solution during business blueprinting should ideally continue uninterrupted throughout the realization phase of an SAP implementation.

Replacing the traditional waterfall approach with a progressive realization approach enables the business and IT teams to work closely together at virtually every step of the solution realization process. This approach reduces the risk of having the end solution misaligned with the strategic and operational design decisions made during business blueprinting.

Progressive realization of the SAP solution starts with the Business Process Manager-based playbacks that were used during business blueprinting, and gradually refines every aspect of the process definition. It continues with the development of detailed decision logic to reduce SAP transaction data entry at run time, reduce SAP configuration, dynamically guide the process, and optimize various kinds of routine inline decision-making.

Of course, traditional realization activities, such as SAP configuration, refinement of the user interface, design and build of custom tables, and even some SAP customization, are performed also. Rather than waiting until user acceptance testing to show the business what has been developed, the progressive realization approach uses regular interactions with the business users to confirm that the implementation decisions made thus far align with what the business needs.

The use of IBM Smarter Process for SAP not only encourages and facilitates the progressive realization approach, it also provides the tools and capabilities necessary to reduce SAP customization and configuration. Through the use of SAP Guided Workflow, easy-to-use SAP technical integration and SAP transaction decomposition techniques, IBM Smarter Process for SAP enables IT developers and SAP functional experts to work interactively with the business as the solution is being developed.

Furthermore, extensions to SAP functionality are much easier to build in the IBM Smarter Process platform, as opposed to SAP tools and traditional application development platforms, such as Java or C#. Many types of customizations, such as the addition of new data fields, integration with non-SAP applications, and extended business rule logic, can be developed in IBM Smarter Process for SAP in a fraction of the time compared to traditional methods.

This combination of iterative solution realization and rapid development techniques presents the entire SAP solution development and deployment team with a tremendous opportunity to help ensure the strategic fit of the solution, while reducing implementation project time, risk, and cost. Time and cost savings of using progressive realization techniques with IBM Smarter Process for SAP can be in the range of 10 - 15%.

Rollout simplification and acceleration

One of the most effective ways that IBM Smarter Process for SAP adds value during the rollout phase is by enabling the process orchestration layer to accommodate slight variants of the core global process, rather than using the traditional SAP global template approach. Orchestrated processes reduce the number of SAP process variants required for large global rollouts, or implementations that span multiple business units or geographical areas.

This is accomplished through the use of decision tables, decision trees, and natural language action statements, all of which can easily be changed by the business. By embedding process variability into a single core definition of a process, the SAP implementation teams can reduce the time, effort, and complexity of typical SAP rollouts while helping to ensure the adoption of consistent business practices across the enterprise.

One of the benefits of orchestrating SAP processes is to simplify the work content of each of the key process steps. Simpler work requires less training, and consequently less training material and material development. SAP Guided Workflow enables process instance data, constants, and data from previous SAP screens used earlier in the process, to be automatically passed into downstream SAP screens.

This ability to intelligently put SAP transactions into the correct *mode*, reduce data entry, and conditionally enter correct screen values, simplifies and reduces work at each process step. It also reduces the number of data fields that a user needs to remember and enter for each specific process scenario.

When used together, SAP Guided Workflow, IBM Business Process Manager SAP technical interface pattern technology, and dynamic coach generation capabilities enable the rapid creation and maintenance of custom SAP user interfaces. They do so without relying upon the traditional costly ABAP and Java approaches.

Creating user interfaces that combine data entry from different SAP transactions or non-SAP systems, reducing the number of fields required, and integrating related lookups directly in the transaction, can dramatically improve both the user experience and business outcomes. Technical integration inserted between SAP transactional steps can also substantially reduce or even eliminate many of the manual steps, both inside of and between transactions, that would ordinarily be required using a traditional SAP implementation approach.

Another key benefit of IBM Smarter Process for SAP during rollout preparation is to reduce and simplify localization activities. Small variations and tweaks to global processes can easily be made using the same rapid design and deployment approach that accelerates business blueprinting and solution realization.

New combinations of business parameters that would ordinarily indicate the need for modifications of the global template can now be incorporated inline in a single global process. The new combinations can be incorporated without extensive changes to the core global process, or a new localized variant of the global process.

Other classic SAP localization configuration activities, such as the setup of local tax tables, new currencies, and unusual units of measure, can benefit from SAP Guided Workflow. Because the SAP transactions used to set up most configuration parameters in SAP are no different technically from most other SAP transactions, the same benefits of passing data into business-focused SAP transactions using SAP Guided Workflow apply equally well to most SAP configuration screens.

Accordingly, the SAP implementation team can quickly set up an SAP Guided Workflow (with screen parameter passing) to semi-automate many of the SAP configuration steps. This reduces implementation time and cost, and improves configuration process consistency.

Lastly, the set of rollout activities that are required for a successful implementation can be modeled, played back, and orchestrated using IBM Business Process Manager. This approach enables the successful coordination and completion of a highly complex set of interrelated rollout activities across organizations that traditionally have suffered from lack of visibility and control.

Being able to see the status of any activity or activity set in real time provides project managers and business stakeholders with the tools necessary to help ensure that key rollout activities remain on track. When orchestrated, these now-proven deployment processes can be used for subsequent rollout activities in other geographical areas or lines of business, with equal or greater benefit.

Global business template value acceleration

A common business scenario for new SAP implementations, instance consolidations, and SAP re-implementations is the use of an SAP global business template to help standardize business processes, policies, and practices. Most often, the SAP global template has been completed long before the SAP platforms to support the new template are ready.

Therefore, a time gap exists between the business definition and the implementation, potentially causing the new global template design to be outdated even before it is implemented. More importantly, it delays the potential business value that the new global process template can bring.

IBM Smarter Process for SAP can help to address both of these issues by providing a dynamic SAP process execution platform that can enable a new global process template to be implemented in whole or part on one or more existing SAP instances.

Because the process layer is decoupled from the application layer, IBM Business Process Manager, often in conjunction with IBM Operational Decision Manager, can dynamically point to different SAP runtime environments during process step execution. This enables a single global process definition to be easily implemented across different SAP instances.

Along with this foundational dynamic execution capability, IBM Smarter Process for SAP enables comprehensive value realization schemes to be implemented directly in the process layer. KPIs, SLAs, and other business metrics can now be calculated and managed outside of the SAP layer for comprehensive coverage across applications, consistency, adaptability, and ease of change.

Remediation processes, or even a more complete operational playbook, are coordinated directly inline with the core business process, helping to further accelerate the business value that the new global template can bring to the business. Perhaps equally important, it enables much of the global process template to be tested independently of a new technology implementation of SAP, to accelerate business value and reduce project and organizational risk.

Value-driven transformation

Most ERP implementations do not meet their business case. The Panorama Consulting 2014 ERP Implementation Success Rate report indicates that the average ERP implementation duration is 16.3 months, 54% go over budget, 72% failed to meet their projected schedule, and 66% realized half or less of hoped-for or promised business benefits. This is not a new phenomenon, because the complexities of ERP system implementation have perplexed businesses and consultancies alike for decades.

Unfortunately for most companies implementing SAP, value management is typically an afterthought. Intensive focus throughout most of the SAP implementation is, almost exclusively, on only those essential activities required for the successful operation of the SAP solution. These essential activities include data cleansing, data migration, core process definition, user training, and operational management.

Although essential, most of these activities are not focused on the kinds of additional optimization activity required to successfully deliver or exceed the SAP business case. The lack of value management focus during the SAP implementation project is clearly one of the key factors in the failure of most SAP implementations to even achieve their base business case.

Value-Driven Transformation (VDT) is an approach for value realization that marries *visibility of important business objectives and measures* with *orchestrated value management processes mapped to organizational capabilities and responsibilities*. An SAP implementation project, either for a net new implementation or a major version upgrade, is the optimal time to introduce VDT techniques. VDT, however, can be introduced at virtually any point in the SAP implementation lifecycle.

The VDT lifecycle starts by agreeing upon and carefully documenting important business objectives. Teams then analyze how the various layers and elements of the business organization contribute to or directly manage relevant aspects of processes related to these objectives. This matrix of the business organization, business objectives, business measures, and contributing processes are mapped into a design framework that IBM calls the *Operational Playbook*.

The Operational Playbook is then reviewed by the organizational leaders responsible for the processes that contribute to the performance of these key business measures. When approved, the operational playbook is then translated into the various layers of the IBM Smarter Process for SAP architecture responsible for gathering, analyzing, and acting upon the business events that drive action.

Figure 4-31 shows examples of key components of the VDT approach described earlier with playbook orchestration implemented as a business process in IBM Business Process Manager, which actively orchestrates all of the activities in the organization required to remediate negative deviations from observed operational metrics and KPIs.

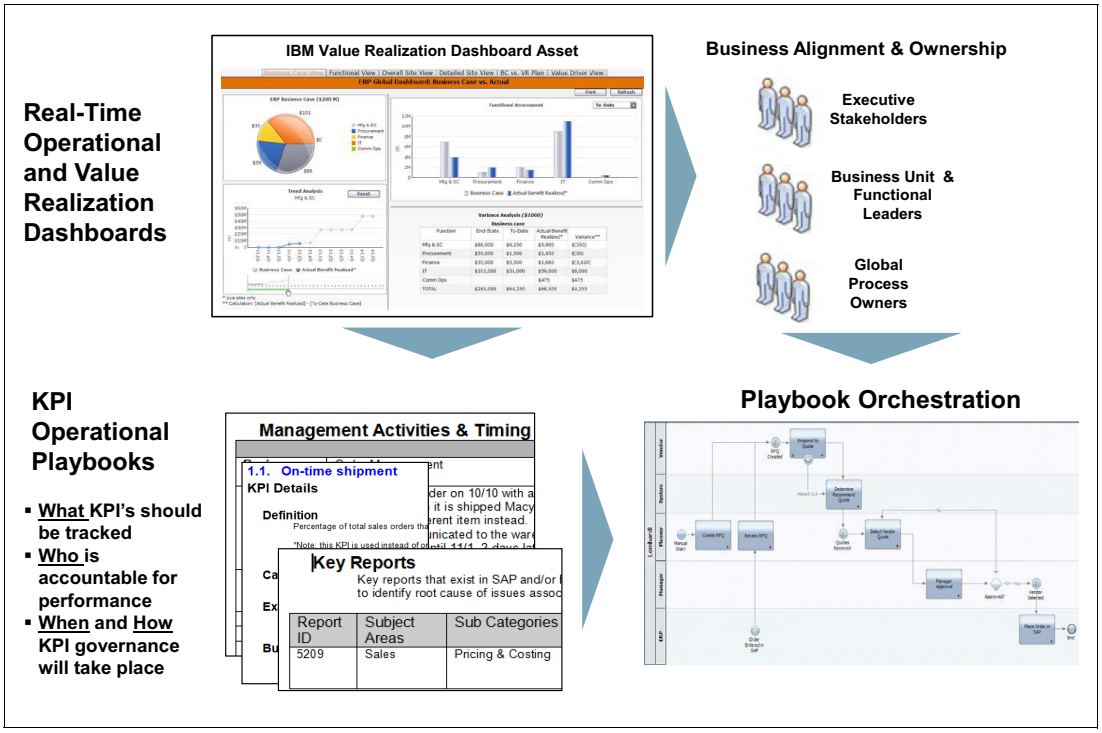

Figure 4-31 IBM Value-Driven Transformation approach

VDT uses the active business performance optimization architecture, introduced in 4.3, “SAP active business performance optimization architecture” on page 78, to gather, correlate, and analyze the various sources of business metrics and events necessary to determine the need for value management action. These sources include SAP applications, non-SAP applications, and various sources of KPIs, KPI trends, and business alerts.

When action is indicated, IBM Business Process Manager then orchestrates all of the activities required to remediate the negative or potentially negative business situation identified by the event processing layer. Should the business fail to run any aspect of the orchestrated playbook, additional escalation and remedial action are then initiated automatically by the orchestrated playbook processes.

This closed-loop system, consisting of event gathering, event correlation, action management, and operational playbook orchestration, delivers a highly reliable solution for improving the value that can be derived from virtually any SAP or heterogeneous business application landscape.

Traditional business intelligence treatments, based on passive analytics, provide much-needed business value around trend-centric optimization opportunities. However, they do little to identify and manage business optimization opportunities in-line with the running business processes.

Ideally, VDT should be integrated into the SAP implementation program from the start of business blueprinting. The careful dissection and analysis of the business process that occurs during blueprinting and detailed process design is the most efficient and effective time to adopt VDT.

The business process walk-throughs, KPI definition exercises, and business process re-engineering activities performed during business blueprinting are the same set of foundational activities required to properly define and build a VDT-based business optimization platform.

IBM Smarter Process for SAP contains all of the technical and business capabilities needed to effectively design, test, and implement a VDT program based on relatively homogeneous SAP, or a heterogeneous application, environment. Event capture, event management, event correlation, KPI definition, business roles, orchestrated value management processes, and performance tracking are all capabilities provided by IBM Smarter Process for SAP.

4.7.2 Post-implementation value augmentation

How IBM Smarter Process for SAP creates value during an SAP implementation project is outlined in 4.7.1, “IBM Smarter Process for SAP in the phases of an SAP project” on page 107.

This section describes the value of IBM Smarter Process for SAP beyond the initial implementation project proper.

Introduction to post-implementation project types

Between the two extremes of net new SAP implementations and major version upgrades of SAP lies an additional spectrum of business value optimization potential. SAP implementations typically struggle with operational visibility, solution flexibility, and business agility for many years after the initial implementation. These business gaps can be easily addressed with IBM Smarter Process for SAP capabilities.

SAP process innovation

One of the easiest ways to get started with IBM Smarter Process for SAP in an existing SAP implementation is to identify one or more business processes that suffer from the typical SAP process optimization gaps described throughout this chapter.

After reasonable process candidates have been identified, a process discovery workshop can be used to clarify the existing challenges with the business process, identify improvement goals, and set about the business of creating a high-level solution and project plan. IBM can assist organizations with every phase of this process, from identifying potential candidates, to conducting the workshop, to developing an effective project plan and approach.

For a mature SAP implementation, starting the IBM Smarter Process for SAP journey with a single process is currently the approach that IBM suggests. A project sufficiently small in scope, typically 90 to 120 days, enables the development of enough functionality to clearly demonstrate to the business the value that IBM Smarter Process for SAP can provide, while keeping the scope small enough to prevent unnecessary risk to the business or to the IBM Smarter Process for SAP journey.

This small pilot project should include as many of the key IBM Smarter Process for SAP capabilities as make sense for the pilot project. Applying a diversity of IBM Smarter Process for SAP capabilities enables the business to gain exposure to, and experience with, a fairly complete set of IBM Smarter Process for SAP capabilities firsthand. The business and pilot project team are now in a better position to determine how to apply these capabilities to subsequent projects.

When the business and IT have practical experience implementing IBM Smarter Process for SAP for one process, an assessment can be made as to which additional processes or improvement opportunities should be addressed next.

Active business performance optimization programs

After experience has been gained with a successful IBM Smarter Process for SAP pilot project, a program-level strategy for optimizing the business should be considered. Typically, these programs include tools and techniques that help select good business optimization candidates, enforce architecture and business practice discipline, and provide the standards and governance necessary for success.

IBM Smarter Process for SAP Center of Excellence

Sometimes programs take the form of a CoE. Staffed by full-time or part-time members, these centers can help deliver the organizational improvements and process improvements required to capitalize on the full potential of IBM Smarter Process for SAP.

Application management services as a business optimization program

An established SAP implementation is generally a gold mine of business optimization potential. The teams who support an SAP implementation after go-live are in a unique position to help the business understand where potential opportunities for substantial value realization might be found. Trouble tickets, repeated areas of confusion, and patterns of request for help often indicate where a systematic approach to addressing a business scenario using IBM Smarter Process for SAP might be beneficial.

4.8 Conclusion

This section summarizes the capabilities and value that IBM Smarter Process for SAP provides to SAP implementations.

4.8.1 Capability and value summary

IBM Smarter Process for SAP provides the following key values:

- Value for the SAP implementation program

  IBM Smarter Process for SAP provides substantial value for every phase of the SAP implementation program. Gross savings of up to 25% are achievable for SAP implementation projects that include substantial configuration or customization. Net savings of up to 15%, after absorbing the costs of IBM Smarter Process software, hardware, and related services, are possible.

- Post-implementation value summary

  The largest savings potential, however, is after the implementation program has been completed. Several benefit illustrations demonstrate that the five-year value of applying IBM Smarter Process for SAP to help manage the business can easily pay for the entire SAP implementation and much more.

4.8.2 IBM Smarter Process for SAP Affinity Analysis and Business Value Assessment Workshop

The IBM Smarter Process for SAP Affinity Analysis and Business Value Assessment tool is an integrated framework that is designed to quickly assess the potential value of IBM Smarter Process for SAP for a part of, or an entire, SAP implementation. The tool is based on a progressive value discovery model that enables IBM teams and qualified IBM Business Partners to assess the qualitative and quantitative benefits of IBM Smarter Process for SAP at virtually any phase of engagement with an SAP client.

SAP, heterogeneous, and non-SAP processes are assessed against a set of weighted process characteristics. Some of these characteristics are inherent to the nature of the process. For example, Order To Cash was, is, and always will be an end-to-end business flow. Business-to-business (B2B) transactions always involve collaboration outside the enterprise. Sarbanes-Oxley related processes are always driven by regulatory requirements, and so on.

Other characteristics, however, are predominantly client-specific. Attributes such as average number of daily instantiations, how mature the organization is in running the process, how many approval paths are required, and so on, vary widely from organization to organization. Other process characteristics are hybrids. This tool and its associated workshop helps organizations to quickly identify the potential value of IBM Smarter Process for SAP for their SAP implementations by providing the following capabilities:

- Quickly identify, quantify, and prioritize IBM Smarter Process for SAP opportunities in the organization.
- Analyze hundreds of processes at a time.
- Adaptable to the organization and approach:
  - Process attributes
  - Business outcomes
  - Capability mapping

- Prioritize and sequence optimization opportunities.
- Build the business case with confidence.

The IBM Smarter Process for SAP Affinity Analysis and Business Value Assessment tool comes prepackaged with a few hundred of the most common SAP processes and their inherent attributes. IBM and IBM Business Partner teams can use the tool to quickly assess the high-level applicability of IBM Smarter Process for SAP to a large number of business processes and scenarios, and then build on earlier assessments as the value discovery progresses downstream. The following list describes the outputs of the tool:

- Overall qualitative value assessment
- Overall quantitative value assessment
- Ranked and scored process affinity list
- IBM Smarter Process for SAP capability matrix

IBM and IBM Business Partner teams typically augment the raw tool output with a sequenced set of project plans and a suggested roadmap for success. The approach facilitated by this tool enables businesses and technical teams to quickly transcend the theoretical value of discussions based on architecture and solutioning to an operational analysis of the actual processes in scope.

This approach helps to determine what IBM Smarter Process for SAP can specifically do to help optimize each business process. It also improves the confidence that the business can have in the value assessment.

Figure 4-32 illustrates how process attributes are mapped, scored, and ranked within the tool.

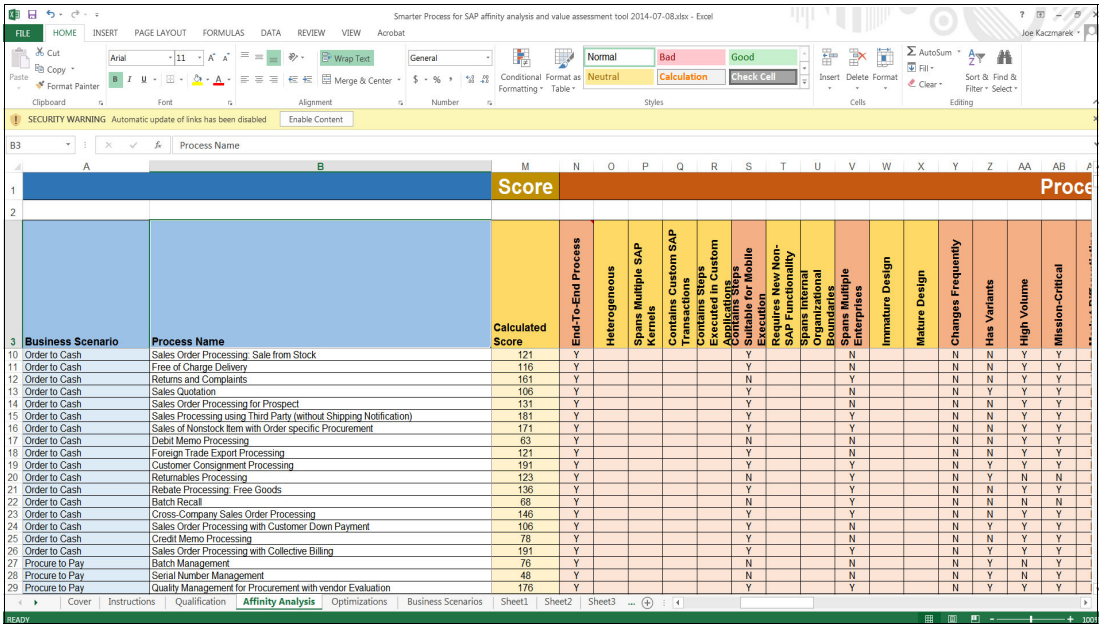

Business Scenario	Process Name	Calculated Score	End-To-End Process	Heterogeneous	Spans Multiple SAP Kernels	Contains Custom SAP Transactions	Contains Steps Executed in Custom Applications	Contains Steps Suitable for Mobile Execution	Requires New Non-SAP Functionality	Spans Internal Organizational Boundaries	Spans Multiple Enterprises	Immature Design	Mature Design	Changes Frequently	Has Variants	High Volume	Mission-Critical
Order to Cash	Sales Order Processing: Sale from Stock	121	Y					Y			N			N	N	Y	Y
Order to Cash	Free of Charge Delivery	116	Y					Y			N			N	N	Y	Y
Order to Cash	Returns and Complaints	161	Y					N			Y			N	N	Y	Y
Order to Cash	Sales Quotation	106	Y					Y			N			N	Y	Y	Y
Order to Cash	Sales Order Processing for Prospect	131	Y					Y			N			N	N	Y	Y
Order to Cash	Sales Processing using Third Party (without Shipping Notification)	181	Y					Y			Y			N	N	Y	Y
Order to Cash	Sales of Nonstock Item with Order specific Procurement	171	Y					Y			Y			N	N	Y	Y
Order to Cash	Debit Memo Processing	63	Y					N			N			N	N	Y	Y
Order to Cash	Foreign Trade Export Processing	121	Y					Y			N			N	N	Y	Y
Order to Cash	Customer Consignment Processing	191	Y					Y			Y			N	Y	Y	Y
Order to Cash	Returnables Processing	123	Y					N			Y			N	Y	N	N
Order to Cash	Rebate Processing: Free Goods	136	Y					Y			Y			N	N	Y	Y
Order to Cash	Batch Recall	68	Y					N			Y			N	N	N	N
Order to Cash	Cross-Company Sales Order Processing	146	Y					Y			Y			N	N	Y	Y
Order to Cash	Sales Order Processing with Customer Down Payment	106	Y					Y			N			N	Y	Y	Y
Order to Cash	Credit Memo Processing	78	Y					Y			N			N	N	Y	Y
Order to Cash	Sales Order Processing with Collective Billing	191	Y					Y			Y			N	Y	Y	Y
Procure to Pay	Batch Management	76	Y					N			N			N	Y	N	Y
Procure to Pay	Serial Number Management	48	Y					N			N			N	Y	N	Y
Procure to Pay	Quality Management for Procurement with vendor Evaluation	176	Y					Y			Y			N	Y	Y	Y

Figure 4-32 Affinity analysis and Value Assessment Tool

4.9 Other IBM Software Group publications, assets, and tools

IBM Software Group has published several white papers and presentations to help you succeed with IBM Smarter Process for SAP. These documents include guidance for preferred practices, process design governance, and implementation approaches. IBM also offers a comprehensive consulting engagement to help organizations build a Smarter Process Center of Excellence to institutionalize the mind-set and practices necessary for their success.

4.10 IBM Global Business Services SAP assets and tools

IBM Global Business Services (GBS) SAP consulting practice is adopting IBM Smarter Process for SAP for various aspects of an SAP implementation. IBM RapidPath for SAP method, a comprehensive SAP implementation project method, has been developed to outline precisely how to apply IBM Smarter Process for SAP capabilities across the entire SAP program lifecycle, including support activities during the post-implementation period.

The IBM GBS SAP practice has also developed a set of SAP implementation deployment tools, mapped closely to the IBM RapidPath for SAP method. These tools take advantage of the implementation work simplification afforded by the reduction and elimination of documentation tasks throughout the SAP implementation lifecycle.

The IBM GBS SAP practice has also modeled the SAP processes from their SAP Express Solution in IBM Business Process Manager. This approach enables GBS SAP practice consultants to accelerate the value of IBM Smarter Process for SAP for consulting engagements, and to continuously enrich these process definitions with insights gained from a vast array of SAP implementation programs.

4.11 References

These websites are also relevant as further information sources:

- Panorama Consulting Solutions 2014 ERP Report

  http://panorama-consulting.com/resource-center/2014-erp-report/
- Integrate SAP Processes with IBM Business Process Manager

  http://www.ibm.com/software/integration/business-process-manager/sap-integration/
- IBM Business Process Manager V8.5.5 documentation

  http://www.ibm.com/support/knowledgecenter/SSFPJS_8.5.5/com.ibm.wbpm.main.doc/kc-homepage-bpm.html
- Business Process Manager product documentation for V8.5.5 (downloadable information center)

  ftp://ftp.software.ibm.com/software/integration/business-process-manager/iehs/

5

Mobile access for SAP

Digital technologies are driving major economic shifts and, with the spread of mobile capabilities, new business challenges arise for each industry. These new technologies are transforming consumer expectations, internal organizational models, and the external competition. In general, organizations feel the pressure to become more consumer and client-centric to stay successful.

Enabling mobile access to SAP business functions and data for enterprise clients, employees, and business partners is a typical requirement for SAP projects.

This chapter describes the enterprise mobile access capabilities, available from IBM for SAP solutions in a heterogeneous enterprise, that provide a foundation and a comprehensive set of resources required to build a system of engagement around SAP. The key principle of this architecture is using best-in-class software for the enterprise mobile platform that is provided by IBM MobileFirst, and separating it from vendor-specific mobile environments.

This chapter includes the following topics:

- 5.1, “IBM MobileFirst overview” on page 118
- 5.2, “Spectrum of mobile app development approaches” on page 119
- 5.3, “IBM MobileFirst for SAP architectures” on page 121
- 5.4, “Optional components driving enhanced features in mobile architectures” on page 136
- 5.5, “Lessons learned from actual projects” on page 140

© Copyright IBM Corp. 2014, 2015. All rights reserved.

5.1 IBM MobileFirst overview

IBM MobileFirst is a comprehensive, market-leading enterprise mobile offering portfolio that enables IBM customers to fully embrace mobile technologies. Figure 5-1 outlines the overall architecture and key capabilities of the IBM MobileFirst Platform.

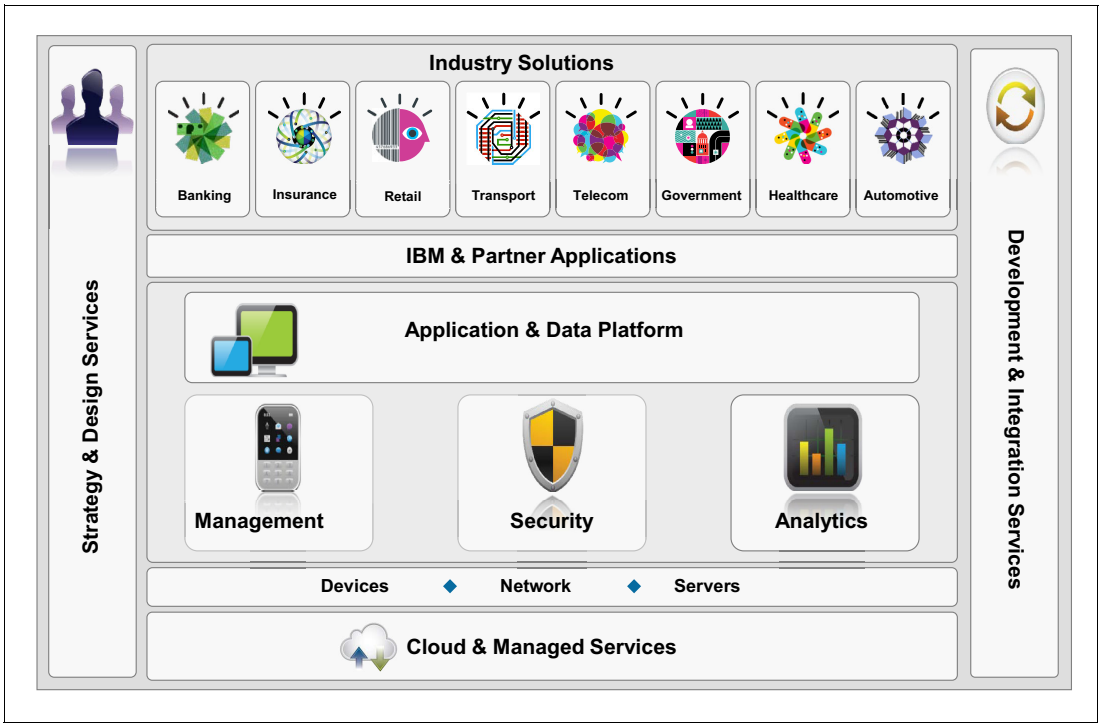

Figure 5-1 IBM MobileFirst architecture

The following list describes the main components of the IBM MobileFirst portfolio:

- Application (app) and data platform. The key capabilities in the platform are oriented to help companies build and deliver engaging mobile solutions more quickly, with higher quality, and at lower cost. Key assets in this space include IBM Worklight® and IBM Rational tools for building and testing mobile assets.
- Management. The need for mobile device and application management, given the growing trend in organizations to enable employee mobility based on the *bring your own devices* (BYOD) principle, is unprecedented. The IBM MobileFirst management capabilities provide a unique solution for all enterprise devices from a single pane of glass, dramatically simplifying the management process.
- Security. Mobile platform security capabilities are critical for mobile enablement, because mobility represents both new opportunities and new threats from a security perspective. The IBM MobileFirst security solutions address the opportunity to make better security decisions based on the context and granularity of access to an application in the mobile context. As an example, many retailers and branch banks want tablet solutions, but they do not want them to work when they are outside the footprint of the store.

IBM MobileFirst security solutions also address a wide set of security threats that are emphasized with mobile enablement. For example, IBM believes that security vulnerability scanning for mobile apps is critical. Therefore IBM added a rich capability of scanning iOS and Android mobile components during the development cycle, to ensure a high code quality level. This capability is especially useful when third parties are involved in building mobile applications that represent a company's brand.

- Analytics. Mobile analytics capabilities are important to ensure a more engaging and higher-quality mobile experience. To achieve this goal, companies need to be able to see what their clients are doing with mobile apps, discover where they are struggling, and where they must wait for too long before taking the next action. Ideally, this discovery should take place before the mobile application is released to an app store, rather than learning that it does not meet client requirements two weeks later.

  The IBM MobileFirst portfolio contains an industry-leading customer experience analytics component. With this component, you can follow (similar to a flight data recorder) all of the swipes, gestures, and screens that the users go through, which rapidly helps teams to build better mobile experiences.

- Strategy and design services. Experienced mobile consultants in IBM consulting services can help clients to explore, assess, and plan mobile solutions, and to prioritize the most important actions to take. Having an efficient design is critical for a mobile solution and, with the IBM Interactive team as part of the consulting organization, IBM is able to help clients to build truly world-class mobility solutions.

- Development and integration services. The IBM MobileFirst offering goes beyond the strategy and design, to help teams build and integrate mobile solutions into the fabric of their business, which is often one of the hardest challenges. IBM consultants have the skills to help organizations build apps from the beginning, in addition to helping them to size and rebuild the infrastructure that they need for successful deployment of the new mobile solutions.

5.2 Spectrum of mobile app development approaches

A spectrum of implementation choices is available for mobile applications in the market. No one perfect answer exists for the choice of implementation for a generic mobile application, and all of the choices across the spectrum have their advantages and disadvantages. Therefore, the challenge for mobile development teams is to understand the trade-offs between the technologies, and make a choice based on the specific application requirements.

IBM MobileFirst supports all variances of these development approaches with the Worklight Platform. Especially in areas where business data is in SAP enterprise resource planning (ERP) systems, it is essential to have a certain kind of flexibility, because the SAP domains provide a diverse interface and technology mix. This mix can range from proprietary up to open standards-driven integration approaches.

Figure 5-2 shows various mobile development methods and their key characteristics.

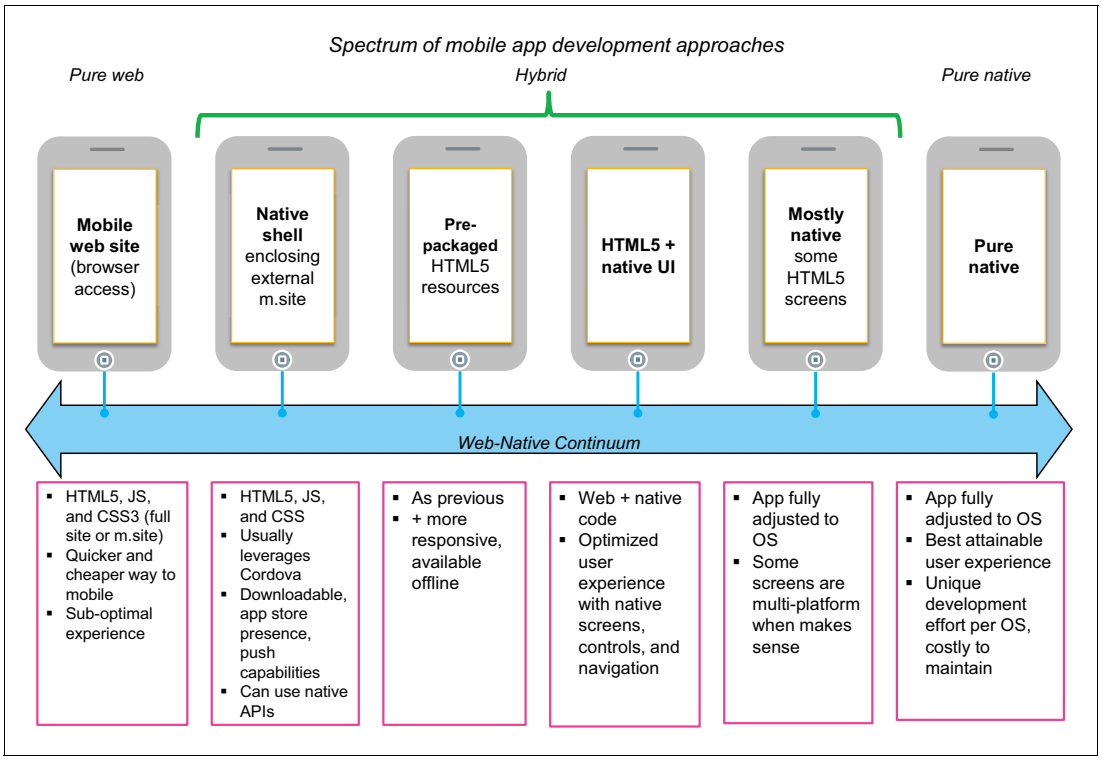

Figure 5-2 Web to native application development continuum

Native application implementation has the advantage of offering the highest fidelity with the mobile device. Because the application programming interfaces (APIs) used are at a low level, and are specific to the device for which the application is dedicated, the application can take full advantage of every feature and service exposed by that device.

Native implementations of mobile apps are completely non-portable to any other mobile operating system. A native Apple iOS app must be totally rewritten if it is to run on an Android device. That makes this choice a costly way of implementing a mobile business application.

A feasible approach is to implement a mobile business app that is a standard *web application*, using responsive web design principles, such as special style sheets to accommodate the mobile form factor and approximate the mobile device “look and feel.” Mobile apps that are implemented with this approach support the widest variety of mobile devices, because web browser support for JavaScript and Hypertext Markup Language revision 5 (HTML5) is fairly consistent.

Several commercial and open source libraries of Web 2.0 widgets can help with this approach. The web programming model for mobile application implementation also has an advantage for enterprises that already have developers trained in the languages and techniques for web application development. The disadvantage of pure web application implementation is that such apps have no access to functions and features that run directly on the mobile device, such as the camera, contact list, and so forth.

Hybrid mobile application implementation is a form of compromise between pure native implementation and pure web implementation. In this approach, developers write the mobile apps using industry-standard web programming languages and techniques, such as HTML5 and JavaScript. However, they package the app into a natively installable format that is distributed through the app store mechanism.

Hybrid apps are linked to additional native libraries that enable the app to have access to native device features from the single application code base within the Worklight family. For example, a hybrid Worklight app includes, ready-to-use, the actual Apache Cordova capabilities that greatly enhance the access to the specific native mobile device functions.

5.3 IBM MobileFirst for SAP architectures

This section describes various common mobile architectures that can be used when integrating SAP-based business modules in a mobile experience required for a system of engagement. An important point to understand is to which extent the organization is adopting the SAP implementation roadmap. Some customers run low-level SAP releases and are hesitant to upgrade their systems regularly.

Alternatively, others are early adopters who upgrade quickly, move the latest SAP releases to production environments, and participate in the SAP Ramp-Up programs. Based on the customer approach to adopting the SAP implementation roadmap, and the existence of system components, different mobile integration patterns can be chosen to accomplish the wanted mobile solution.

This section highlights common (managed) mobile architectures that provide the capabilities of a mobile enterprise application platform (MEAP). The introduction of a powerful MEAP enables enterprises to handle diverse mobile scenarios efficiently.

It is not about getting one mobile application into production, but about meeting the following requirements:

- Driving various mobile solutions across multiple lines of business
- Serving internal and external business users
- Employing the solutions on multiple mobile operating systems
- Functioning on different mobile devices

All of these requirements must be implemented in a secure manner, following corporate security guidelines.

5.3.1 Architecture goals for SAP mobile enablement in a heterogeneous enterprise

Enterprise mobility enablement requires a platform that provides a wide set of functions:

- Rapidly enable access to enterprise applications (SAP and non-SAP), and provide connectivity and data access approaches for mobile applications to interact with the enterprise applications.
- Provide remote device management and mobile device security in line with the corporate security infrastructure.
- Enable effective distribution mechanisms for private enterprise mobile applications.
- Reuse existing or pre-built assets for mobile devices, for example, pre-built mobile applications provided by SAP.
- Exploit unique abilities of mobile devices to provide differentiating business capabilities for companies.

No single facet within the mobile environment defines the system. In defining the mobile platform architecture, the entire mobile landscape must be considered.

Provide enterprise integration

The enterprise mobility platform should be designed as an enterprise-wide asset that provides a common mobile enablement framework for both SAP and non-SAP enterprise applications. Non-SAP enterprise applications can include an organization's internally developed applications (IBM CICS®, IBM IMS™, Java Platform, Enterprise Edition, and Microsoft .Net systems), cloud services, and other third-party commercial products.

Enterprise mobile access to SAP for a heterogeneous enterprise reuses and extends existing enterprise integration assets, and integrates with existing enterprise security and endpoint management infrastructure.

Enable multiple mobile device platforms and mobile application types

Enterprise mobile access to SAP must support multiple mobile device platforms in an economical way, and a spectrum of implementation choices for mobile applications. Based on business requirements, mobile applications for SAP can be any of the following types:

- High fidelity native mobile applications to provide differentiating mobile experience to clients.
- Hybrid mobile applications with varying degrees of combined native and embedded web components. The hybrid model enables you to use a common code base to target different mobile platforms, while enabling access to native capabilities of mobile devices, such as camera, location, and so on.
- Pure web-based mobile applications based on responsive web design for entry level mobile access to SAP and non-SAP data.

All of the choices across the spectrum have their advantages and disadvantages.

Support prepackaged SAP mobile applications and domain decoupling

SAP and its partners provide a set of prepackaged mobile applications. Deployment of such mobile applications enables acceleration for mobile access enablement when a business function differentiation is not a core requirement. This can be the case when providing mobile access for company employees to internal business support systems.

The key consideration in planning adoption of prepackaged SAP mobile applications is the existence of requirements to change ready-to-use application functions to reflect enterprise-specific needs. Such changes to ready-to-use applications can be further classified as *configuration changes* (no application code change required) and *customization* (application code changes are required).

Prepackaged mobile applications can provide a substantial business value, when such applications meet business requirements *as is*, without customizations, or only through well-defined configuration changes supported by SAP. With prepackaged applications, the issues of trust, login, user management, upgrades, and system access are solved with minimal effort.

The long-term return on investment (ROI) on SAP implementations is, in large part, directly linked to the level of SAP enhancements and modifications. Extensive SAP customizations lead to a dramatic increase in SAP upgrade costs at a future date. The problem is retrofitting the customizations made to the prepackaged application when the underlying platform is upgraded to a new major functional release that includes significant changes.

A migration to a new release of SAP business suite will, invariably, be faced with a problem of refactoring earlier enhancements and modifications onto a new platform version, which can become even more costly because of the unpredictable nature of changes in future versions.

If prepackaged applications are a good fit, the enterprise mobile access solution must enable the use of pre-built SAP mobile applications. From an architectural perspective, doing so requires you to maintain two categories of mobile apps in an overall mobile solution for the heterogeneous enterprise:

- Standard mobile platform domain

  An enterprise standard for all custom mobile development, including both SAP and non-SAP enterprise assets.

- SAP mobile platform domain

  Needed to deploy prepackaged SAP mobile applications. This domain should not be used for any custom development, and is treated as a *black box* (only its external behavior is considered). If an SAP prepackaged application requires extensive enhancements and modifications to meet requirements, the application function should be developed on a standard mobile platform.

A key architecture goal is to keep these two categories separated, and ensure that changes in one sub-area do not affect the other one. This goal can be achieved by applying traditional middleware leading practices, which are still valid for mobile solutions.

Provide business differentiation based on the mobile experience

For any user-facing application, a key requirement is that it should be easy and efficient for the target user audience to use it. Usability is even more critical for mobile apps, especially in the business-to-consumer (B2C) environment, where consumers decide intuitively within the first interactions with the mobile app if they like it.

As companies aggressively extend their business presence to the mobile devices of their customers, the ability to differentiate their business services by using *mobile moments* with unique capabilities of mobile devices becomes one of the key requirements for the enterprise mobility platform. Such differentiation requires an ability to provide custom mobile access enablement for enterprise applications.

The enterprise mobility platform, therefore, must effectively support development of custom mobile content for both SAP and non-SAP enterprise applications. The platform must also provide enterprise connectivity and a data access framework for mobile device applications to interact with back-end enterprise applications.

5.3.2 IBM MobileFirst for SAP architecture overview

The enterprise mobile platform architecture shown on Figure 5-3 supports all of the architectural goals outlined earlier in this chapter.

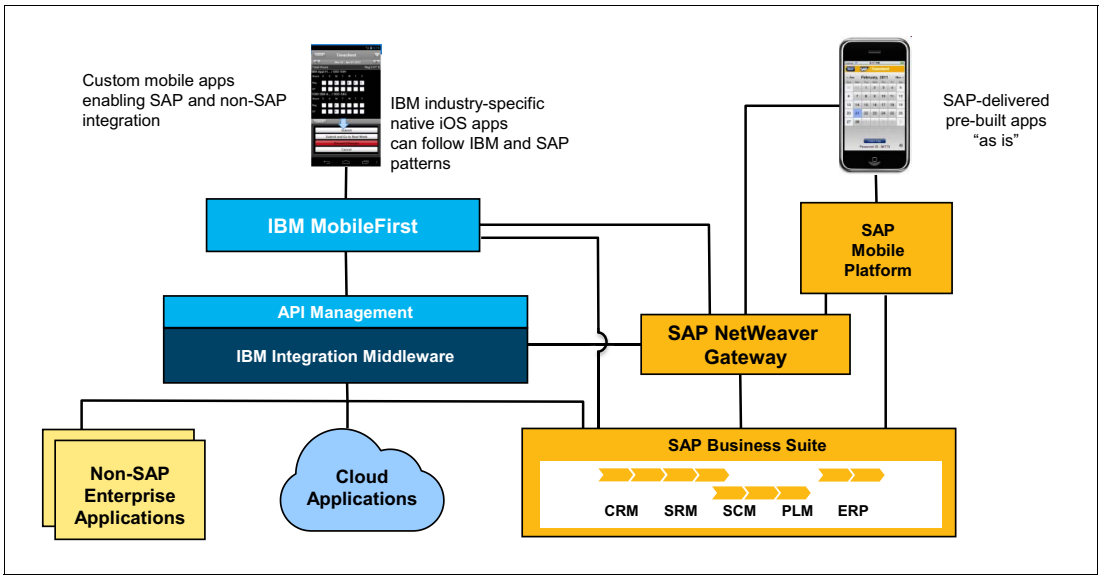

Figure 5-3 Heterogeneous mobile enablement for SAP

The architecture shown in Figure 5-4 on page 126 consists of two technology domains:

- The MEAP domain for all SAP and non-SAP enterprise mobile applications based on IBM MobileFirst.
- The SAP mobile domain treated as a black box and used exclusively for deploying pre-built SAP mobile applications that meet business requirements as is, or only through well-defined and SAP-supported configurations.

IBM MobileFirst is a best-of-class enterprise mobility platform built on open standards and designed for heterogeneous environments, both SAP and non-SAP back-ends.

Besides enabling access to SAP data through SAP integration capabilities that are provided by IBM, MobileFirst provides additional value. An essential differentiator is the horizontal concept of the IBM MobileFirst portfolio, which provides generic mobile enablement frameworks, and is fully capable of supporting any type of system of record rather than favoring a particular system or technology.

The Worklight Platform addresses this need by introducing an architecture that is divided into components. A typical mobile solution consists of a *client component*, the application the user runs on the mobile device. This client component interacts with the central Worklight *server component*. The Worklight server provides enhanced enterprise capabilities, such as a generic security framework, efficient session handling, and enhanced analytics capabilities.

The Worklight server additionally provides components that connect various enterprise data sources to the mobile solution. The Worklight adapter can access mobile-ready interfaces directly from SAP software, or use IBM integration middleware components that provide middleware features to the overall mobile solution.

IBM WebSphere Cast Iron Cloud Integration, IBM Integration Bus, IBM WebSphere DataPower, IBM Business Process Manager, and IBM Operational Decision Manager are typical products that enable such integration capabilities. However, non IBM offerings can also be considered, such as SAP NetWeaver Gateway or SAP Process Integration.

IBM API Management provides companies with the tools for creating, proxying, assembling, securing, scaling, and socializing web APIs. Web API is a server-side programmatic interface to a defined request-response message system, typically expressed in JavaScript Object Notation (JSON) or Extensible Markup Language (XML). This interface is exposed across the web, and is most commonly used for developing mobile applications.

Equipped with a customizable developer portal, IBM API Management enables organizations to attract and engage with mobile application developers to foster an increased usage of the published APIs. The robust administration portal in IBM API Management enables companies to easily establish policies for critical API attributes, such as self-registration, quotas, key management, and security policies. The robust analytics engine provides valuable role-based insight for API owners, solution administrators, and application developers.

The following sections further describe patterns used for integration of SAP business data in an IBM MobileFirst architecture.

- 5.3.3, “Fast-track SAP mobile enablement with IBM Worklight and SAP NetWeaver Gateway” on page 125
- 5.3.4, “IBM MobileFirst integration with SAP with no moving parts” on page 129
- 5.3.5, “Accelerated mobile integration with SAP using IBM WebSphere Cast Iron” on page 129
- 5.3.6, “Full featured mobile integration with SAP using IBM Integration Bus” on page 132

The decision to select components is heavily driven by the specific customer environment and the planned mobile solution. The three steps are as follows:

1. Conduct an assessment of the existing interfaces, APIs, integration capabilities, and predefined governance rules for the customer. This assessment delivers specific fixed points, which are set on the mobile enablement architecture and are not negotiable.
2. Collect critical requirements of the wanted mobile solution. Such requirements can be a need for offline capabilities, specific security and performance aspects, back-end integration requirements, or which mobile devices must be supported.
3. Identify gaps in the existing enterprise architecture and missing infrastructure components.

The consolidated outcome of these three steps gives a good indication of which mobile pattern is most suitable for the specific mobile enablement scenario.

5.3.3 Fast-track SAP mobile enablement with IBM Worklight and SAP NetWeaver Gateway

This integration architecture is based on a relatively new type of SAP interface exposed by SAP NetWeaver Gateway, a recent addition to the SAP integration stack. SAP NetWeaver Gateway enables access to SAP data and functional services through a set of Representational State Transfer (REST) services using the OData and Atom protocols. This interface simplifies access to SAP business system from non-SAP systems of engagement, such as mobile applications, social media, and IBM Collaboration Solutions.

To implement this pattern, the mobile application developer does not need to have deep SAP skills because, from the mobile application perspective, the SAP environment can be seen as just another service provider, a REST API.

However, additional and potentially complex SAP configuration might be required in SAP NetWeaver Gateway to service-enable access to SAP business data. The fact that SAP NetWeaver Gateway is based on Advanced Business Application Programming (ABAP) implies that it is also operated by the SAP operations team within an organization, because the administration of this component requires deep SAP-specific skills.

Worklight fully supports this SAP interface, and provides pre-integrated capabilities to automatically discover SAP services exposed by SAP NetWeaver Gateway, and generate integration code adapters and mobile application templates. This pattern is completely contained within Worklight. It can be used best for fast-track development projects that can produce fully functional application code quickly, and be deployed for quick-win mobile initiatives.

The biggest advantage of this pattern is seen when the specific integration scenario can be served by SAP NetWeaver Gateway services that are included by SAP with the standard delivery. In this case, almost no effort is made to implement the integration, because the interfaces can be used *as is*.

When a customer has the SAP NetWeaver Gateway technology in place, a worthwhile step is to verify whether this component can provide the required interfaces into the SAP domain in a consumable manner. This is also a good starting point from the governance perspective, because the control of these interfaces is still within the SAP domain.

Figure 5-4 describes how Worklight-based mobile solutions can communicate with the REST and OData interfaces provided by SAP NetWeaver Gateway.

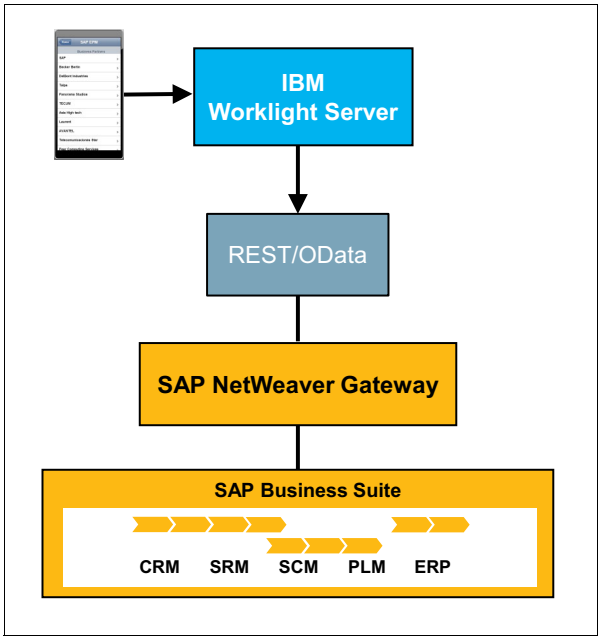

Figure 5-4 Integration using SAP NetWeaver Gateway

Worklight provides a ready-to-use adapter for SAP NetWeaver Gateway. This capability includes a wizard-driven code generation feature to perform dynamic introspection on an SAP NetWeaver Gateway instance, and to create the required artifacts automatically. The generated objects are packaged and deployed automatically within the Worklight environment.

During run time, the Worklight adapter runs the call to SAP NetWeaver Gateway accordingly, and the details of such calls are encapsulated behind a common Worklight adapter programming model that is not SAP-specific. The manual coding effort for the mobile app developer is small in this scenario.

Figure 5-5 illustrates the interactions of underlying components during the design, development, and runtime phases of this scenario.

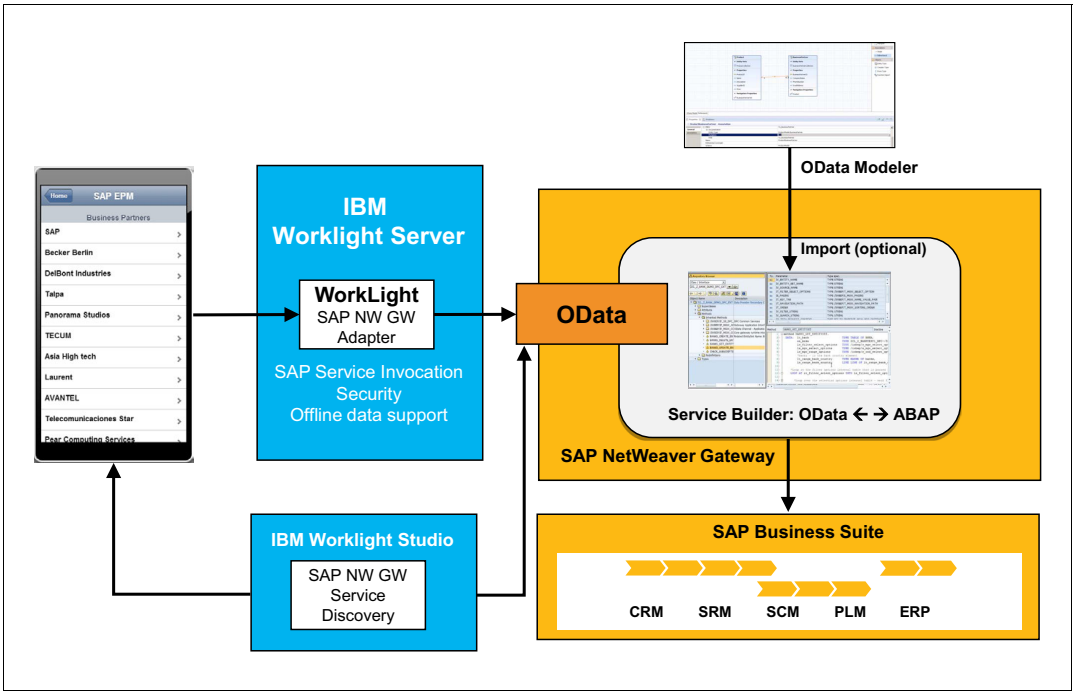

Figure 5-5 Worklight and SAP NetWeaver Gateway

In mobile development scenarios, the development platform must include efficient code generation capabilities to enable the developers to produce high-quality code that encapsulates complex interactions with back-end systems. New mobile apps or upgrades to existing mobile apps must be developed easily and quickly.

Such code generation capabilities are provided by the Worklight adapter for SAP NetWeaver Gateway. The developer can browse the existing OData objects catalog dynamically, and generate the required runtime objects automatically.

Figure 5-6 shows an example of the embedded wizard, which reads the metadata from SAP NetWeaver Gateway and displays a list of available REST and OData services to the mobile developer.

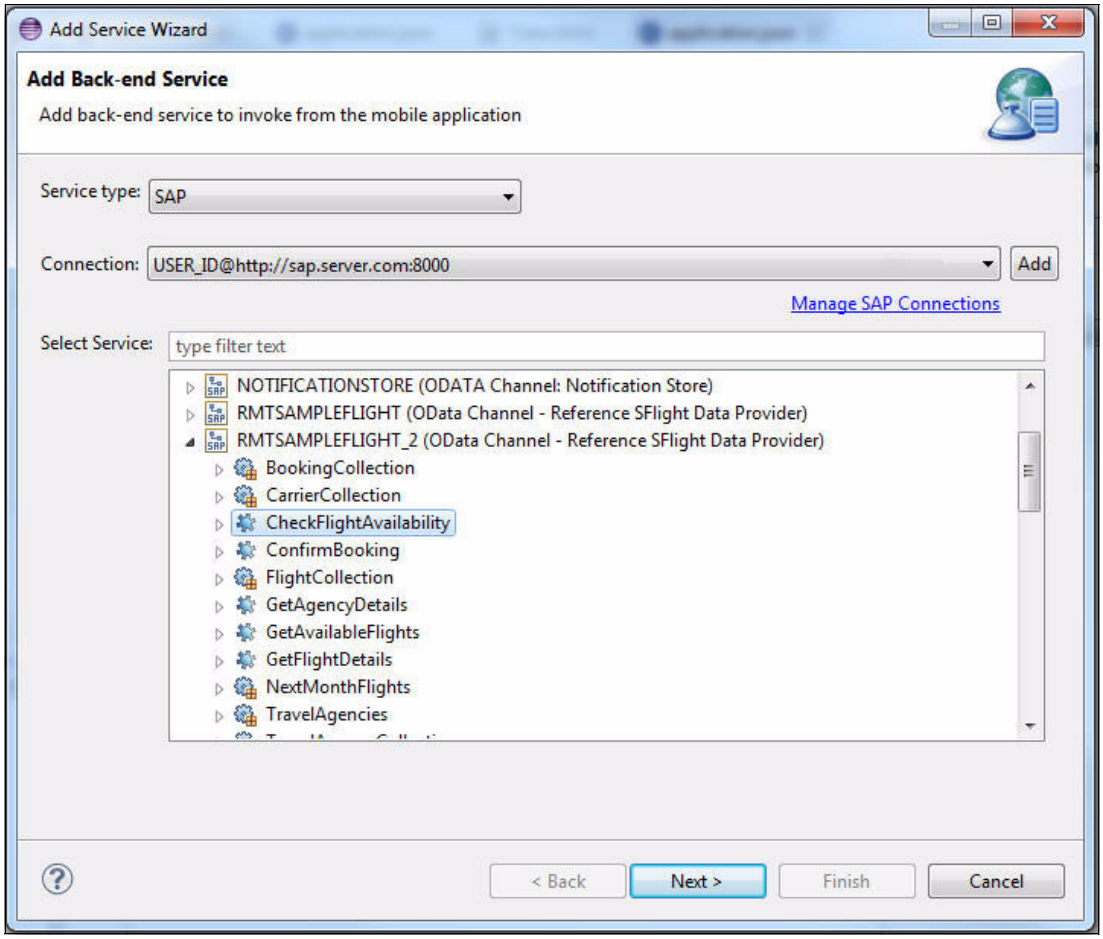

Figure 5-6 Worklight Adapter for SAP NetWeaver Gateway wizard

This capability prevents manual errors by the developer, improves code quality, and increases implementation speed.

Although this integration approach masks the SAP software details from the Worklight developer, it relies on SAP NetWeaver Gateway, and therefore requires the skills of the SAP support team in the organization to configure it. SAP positions SAP NetWeaver Gateway as the strategic integration point for scenarios where users interact with the SAP system through various components, such as fat clients, web browsers, and mobile devices.

These components are the consumers of the SAP business data. For all of these components, the preferred data formats and interaction style with the SAP domain are through OData streams, REST style, and JSON formats. Therefore, the Worklight SAP service discovery wizard will be of tremendous value to SAP customers in the years to come.

5.3.4 IBM MobileFirst integration with SAP with no moving parts

This integration architecture does not require any additional IBM or SAP middleware software. As shown in Figure 5-7, this architecture is based on a direct integration between Worklight server and SAP Business Suite, using the well-established traditional SAP BAPI/RFC interface enabled by SAP Java connector (SAP JCo).

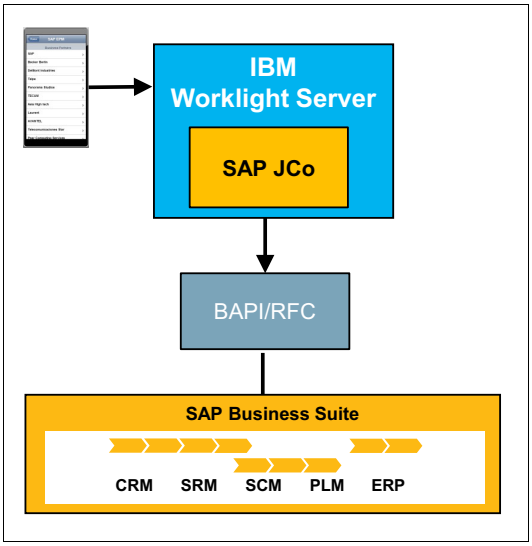

Figure 5-7 IBM MobileFirst integration with SAP with no moving parts

In this case, integration with SAP is implemented using Java integration code, developed based on the SAP JCo interface specification. Subsequently, such Java code is wrapped in an Worklight adapter and is deployed on the Worklight server.

This integration approach enables a direct communication path between Worklight and SAP Business Suite, for example, SAP ERP or SAP customer relationship management (CRM) applications. This integration approach does not require SAP NetWeaver Gateway or any additional IBM middleware software.

The Worklight integration with SAP uses a set of pre-built SAP integration points (BAPIs), which are well-established and documented. Therefore, it can be used with SAP solutions without modification.

From the mobile application developer perspective, this integration option is no different from other options, because the SAP system is exposed to the mobile app using standard Worklight adapter APIs. No additional mobile developer skills are required in this approach. However, on the server side, this approach requires developing custom Java code at the SAP JCo level, and it needs more experienced Java developer skills.

5.3.5 Accelerated mobile integration with SAP using IBM WebSphere Cast Iron

This integration architecture is appropriate when no consumable REST APIs are provided by the SAP domain itself, for example by SAP NetWeaver Gateway. In such cases, IBM MobileFirst enables you to easily include additional IBM middleware components for SAP integration. Cast Iron is a lightweight, efficient, and flexible integration component that can integrate SAP and non-SAP business data.

The main advantage of Cast Iron is simplicity. The complexity of application integration is effectively eliminated because of a wizard-based configuration, rather than a coding

approach. Cast Iron is designed to dramatically accelerate the integration effort required for SAP integration. Cast Iron integration is based on using a rich set of pre-built integration templates, which quickens the integration effort tremendously compared to custom coding.

Integration based on Cast Iron is greatly accelerated, because it includes a rich set of connectors to integrate with nearly any source system. One of these connectors is the IBM WebSphere Cast Iron SAP connector, which enables a two-way communication between Cast Iron and the SAP instance. The connector supports BAPI, RFC, and Intermediate Document (IDoc) interfaces.

These proprietary SAP protocols are the most commonly used SAP integration interfaces, and they enable access to nearly any kind of business data, which is in an SAP (ERP) system. The Cast Iron SAP connector supports any SAP R/3 system, which is based on the ABAP application server (3.1H or later versions) and uses as low-level API the SAP Java Connector (SAP JCo 3.0.x or later). This component uses traditional SAP integration interfaces based on RFC.

A typical SAP implementation exposes a large number of RFC-based interfaces, which are well documented as BAPIs. These interfaces are standard in SAP, and typically do not require any special or additional SAP configuration. This approach differs from the one described in 5.3.3, “Fast-track SAP mobile enablement with IBM Worklight and SAP NetWeaver Gateway” on page 125, where the SAP NetWeaver Gateway configuration is required to expose the interface.

The Cast Iron SAP connector also supports the most popular interaction patterns with the SAP system from a security perspective using a technical user identity, a named user identity, or SAP LoginTickets.

The Worklight Platform contains a pre-built adapter for Cast Iron. This adapter enables the mobile application developer to easily incorporate any available Cast Iron orchestration.

Figure 5-8 highlights how the Cast Iron technology can be used to enable the mobile application to access SAP back-end components.

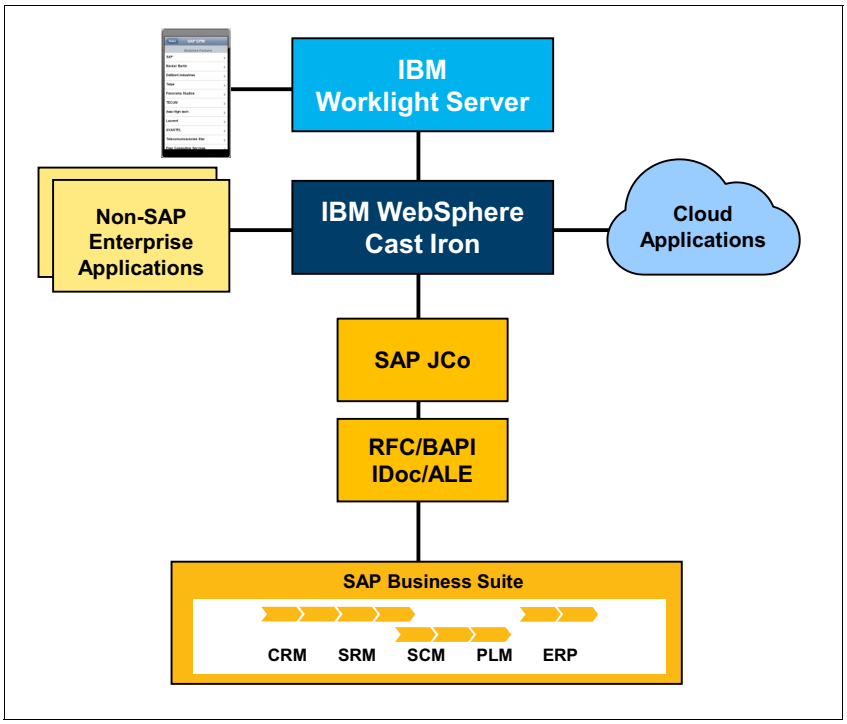

Figure 5-8 Integration using IBM WebSphere CastIron

Cast Iron provides flexible deployment options, because it is available in several form factors:

- IBM WebSphere Cast Iron Hypervisor Edition. A virtual appliance that can be installed on existing servers through virtualization technology.
- IBM WebSphere Cast Iron Live. A multi-tenant, cloud-based platform for integrating cloud and on-premises applications and enterprise systems in a hybrid environment.
- IBM WebSphere DataPower Cast Iron Appliance XH40. A self-contained, physical appliance that provides what is needed to connect cloud and on-premises applications.

Cast Iron offers the developer a client component called IBM WebSphere Cast Iron Studio. This client component enables the generation of integration content regardless of the target form factor. The mobile integration content is built once, and can be deployed on any of the different Cast Iron form factors available.

A unique characteristic of this pattern is the different form factors that Cast Iron provides. It can be run as a traditional on-premises component, or as a cloud-based offering. Besides the choice of form factors, Cast Iron includes more than 75 different connectors to a wide range of popular back-end systems. Therefore, it serves as a complete integration layer.

Figure 5-9 depicts the various components in the integration architecture.

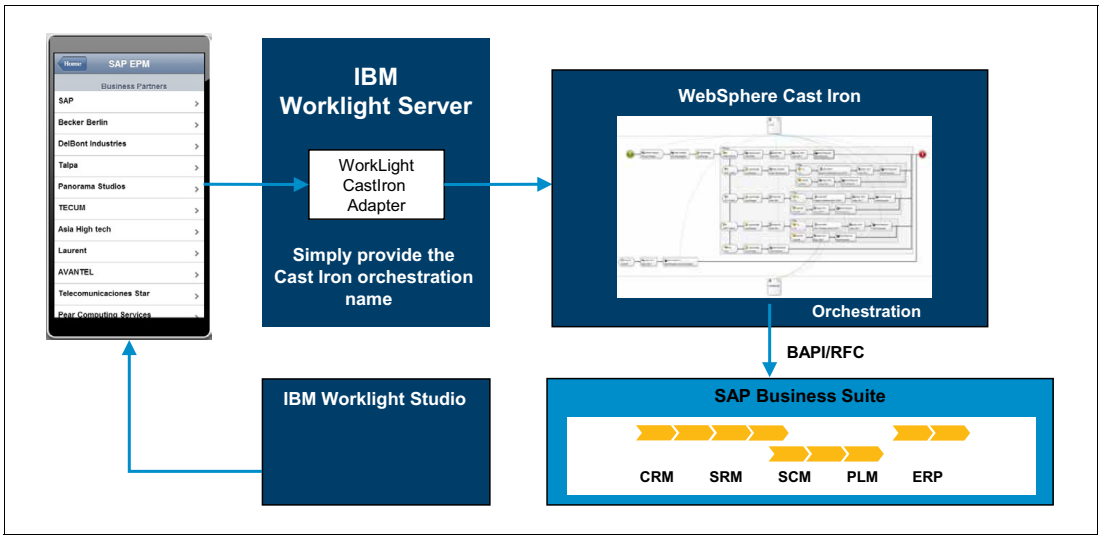

Figure 5-9 IBM WebSphere Cast Iron orchestration

A special feature is the enhanced support for the mobile application integration developers to generate all of the required business objects automatically from the target SAP system. The developer needs to know only the name of the function module or IDoc object; the generation wizard performs the object creation dynamically by reading the metadata from the SAP instance. This enables the mobile application integration developer to produce high-quality code in a short time.

5.3.6 Full featured mobile integration with SAP using IBM Integration Bus

IBM Integration Bus (formerly known as IBM WebSphere Message Broker) is a market-leading middleware component that is used to integrate heterogeneous enterprise systems. The IBM Integration Bus has these features, among others:

- Handling structured, unstructured and binary data
- Synchronous and asynchronous interactions
- Running stateless and stateful integrations
- Routing open standards based and proprietary protocols

IBM Integration Bus provides a variety of options for implementing a universal integration foundation based on an enterprise service bus (ESB). Implementations help to enable connectivity and transformation in heterogeneous information technology (IT) environments for businesses. This integration is critical for organizations deploying service-oriented architecture (SOA), IBM Business Process Manager, existing application modernization, mobile solutions, and third-party products.

For SAP integration, IBM Integration Bus uses the IBM WebSphere Adapter for SAP Software, which employs the most commonly used SAP interfaces, such as BAPI, RFC, Application Link Enabling (ALE), and IDocs. These interfaces can be used in both directions, inbound and outbound, to send business data into SAP environments or to receive business data from SAP environments.

Figure 5-10 highlights how IBM Integration Bus can be used to integrate a mobile solution based on Worklight with SAP.

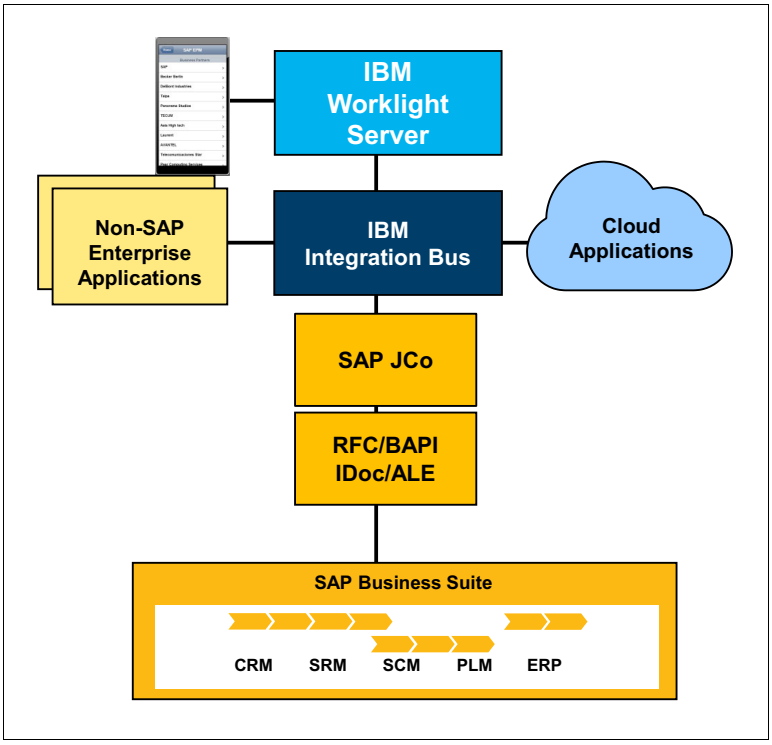

Figure 5-10 Worklight mobile application Integration with IBM Integration Bus

An important point to understand is that the selection of this architecture pattern is driven by either of the following benefits, rather than simply introducing IBM Integration Bus as part of the mobile solution:

- Reusing an existing IBM Integration Bus implemented in the organization
- Introducing an enterprise integration platform in addition to SAP

Enterprises of a certain size typically have a middleware architecture defined at the corporate level; reusing this middleware is a good approach if it is able to provide the interfaces in a suitable format and protocol.

IBM Integration Bus has been used for a long time as a system-to-system integration bus for SAP enterprise applications, and it also has rich capabilities for providing efficient, REST-based integration services. A typical usage scenario is to encapsulate existing and robust IBM Integration Bus message flows and make them available to mobile solutions.

Reusing a common integration governance model across all planned mobile projects is beneficial for the overall mobile experience. Aspects such as an efficient identity mapping across the mobile and the back-end domains, in addition to specific response time requirements for the mobile solution, can be successfully addressed by an IBM Integration Bus middleware layer.

The power of IBM Integration Bus for integration of SAP systems into mobile apps lies in the acceleration of the integration development using patterns. IBM Integration Bus patterns encapsulate the leading practices for mobilizing integration flows and code generation for the Worklight adapter. Using patterns, organizations can achieve a faster time-to-market and better ROI.

Figure 5-11 represents the architectural overview of the IBM Integration Bus Mobile Service pattern.

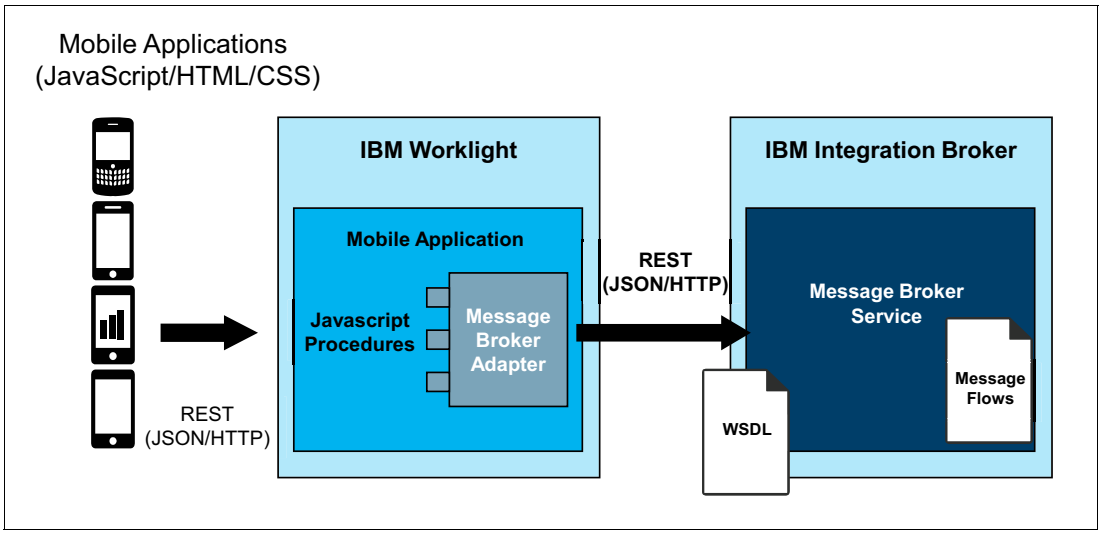

Figure 5-11 Built-in patterns to integrate Worklight and IBM Integration Bus

5.3.7 Access to existing SAP Fiori Apps using IBM MaaS360

SAP introduced the SAP Fiori initiative in 2013, which applies modern user interface (UI) design principles to the existing SAP business modules delivered with the SAP Business Suite. SAP Fiori follows current UI design paradigms to be responsive, personalized, and simple to use.

Additionally, it is consumable by various devices, providing a modernized user experience across the diverse SAP business module landscape.

The applications provided by SAP Fiori can be categorized into several core types:

- Transactional apps. These applications enable business users to perform transactional tasks with the SAP business system.
- Analytical apps. These applications provide the business user role-based insight into key performance indicators of the SAP business system.
- Fact sheets app. These applications are simplified, read-only components displaying contextual information and key facts about core SAP business objects.

SAP Fiori uses core SAP NetWeaver components, such as the SAP NetWeaver Gateway and the SAPUI5 user interface technology. SAPUI5 was an SAP proprietary enhancement of the jQuery framework, and has recently been made open source by SAP with a new synonymous name OpenUI5 (using the Apache License 2.0).

OpenUI5 is based on JavaScript, and supports client-side application rendering using standard browsers on various devices using the HTML5 and Cascading Style Sheets Level 3 (CSS3) standards. From the data binding perspective, OpenUI5 supports different models, such as XML, JSON, and OData.

At the beginning of 2014, SAP claimed to provide a large number of Fiori-based applications, which are split into the core types listed previously. SAP plans to make more Fiori-based applications available in the future, especially in the business-to-employee (B2E) domain.

Therefore, the IBM MobileFirst initiative must be able to provide options that enable developers to use this asset in an easy, seamless, and nondisruptive way.

From a technical perspective, the SAP Fiori applications can be treated like any other web-based intranet application, because the deployment of SAP Fiori applications follows intranet web application design principles.

The IBM MobileFirst portfolio includes IBM MaaS360®, a powerful product set supporting different areas, such as cloud-based mobile device management, mobile application management (MAM), and secure containerization. This support gives organizations the building blocks to separate personal apps data and content from enterprise apps data and content on mobile devices.

Figure 5-12 shows the different MaaS360 usage domains.

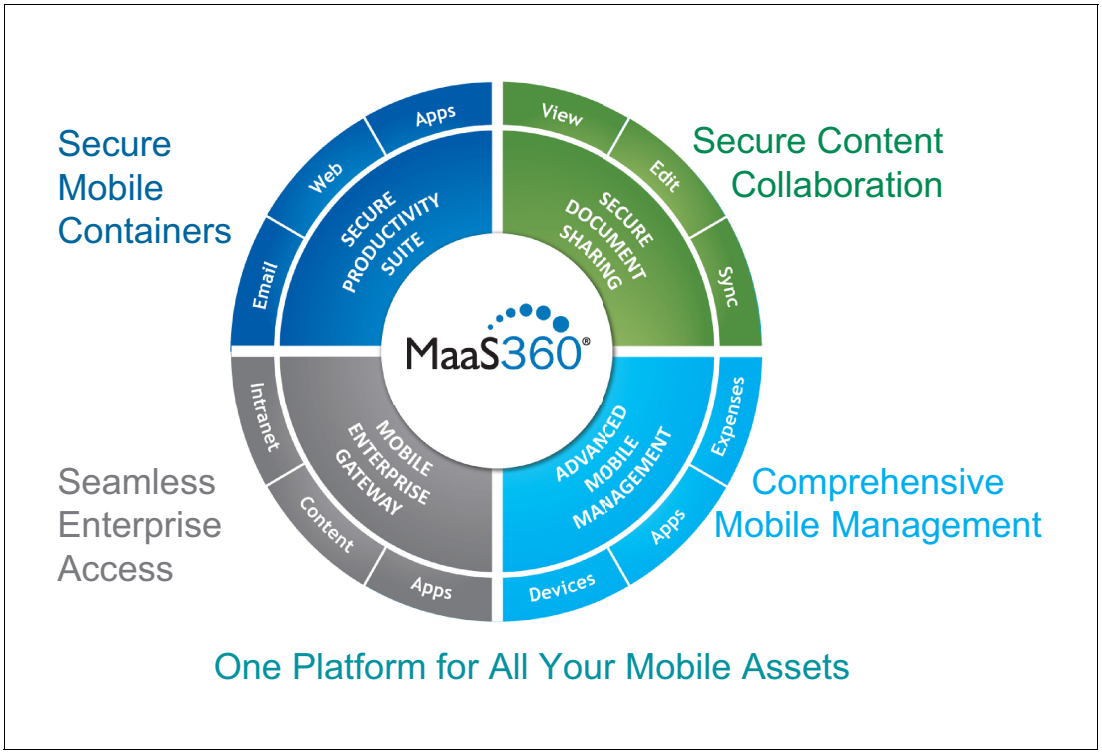

Figure 5-12 MaaS360 overview

For the architecture pattern designed to integrate SAP Fiori mobile apps through MaaS360, the left half of the circle (Secure Productivity Suite™ and Mobile Enterprise Gateway) shown in Figure 5-12 is of most relevance.

The IBM MaaS360 Secure Productivity Suite provides a component called MaaS360 Secure Browser that enables secure access to intranet sites and web applications such as the SAP Fiori apps. MaaS360 Secure Browser enables a fine granular Uniform Resource Locator (URL) filtering mechanism with enhanced features such as restricting cookies, downloads, copy and paste, and printing.

The network component on the provider side is the IBM MaaS360 Mobile Enterprise Gateway, which controls the access of the MaaS360 Secure Browser to internal resources. This approach does not require enabling virtual private network (VPN) on the device.

Figure 5-13 shows the high-level architecture of MaaS360 with SAP Fiori Apps.

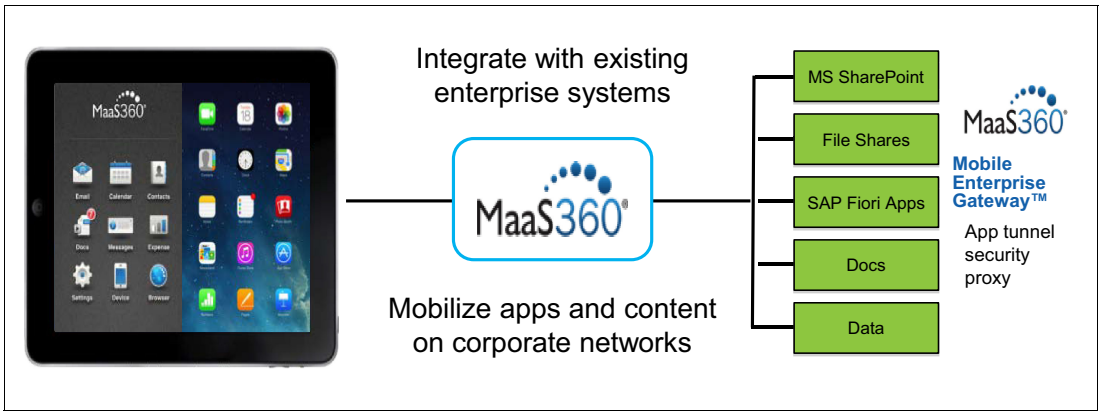

Figure 5-13 Enabling corporate apps for external use

Using MaaS360 Secure Productivity Suite and MaaS360 Mobile Enterprise Gateway with ready-to-use SAP Fiori apps can save a huge amount of development time, if the SAP Fiori apps fit the specific business requirements. With this integration approach, existing assets can be re-used rather than reproduced on new platforms.

5.4 Optional components driving enhanced features in mobile architectures

Based on the mobile scenario and the customer requirements, adding optional components into the mobile architecture to use specific features of those components makes sense. This section describes optional IBM products and offerings that can be added to a mobile architecture for SAP solutions.

5.4.1 Enhancing mobile architectures by adding IBM API Management capabilities

IBM API Management provides companies with the tools for creating, proxying, assembling, securing, scaling, and socializing web APIs. APIs can be defined as programmatic interfaces that expose corporate data assets, such as products, prices, and availability, and that are used by mobile components to interact with the business domain.

Equipped with a customizable developer portal, IBM API Management enables companies to attract and engage with application developers to foster an increased usage of the published APIs. IBM API Management's robust administration portal enables companies to easily establish policies for critical API attributes, such as self-registration, quotas, key management, and security policies. The robust analytics engine provides valuable role-based insight for API owners, solution administrators, and application developers.

IBM API Management is a unified solution, with an intuitive user experience, for managing the complete API lifecycle, from creation, publishing, and adoption to support and monitoring, enabling companies to realize the maximum value from their APIs.

IBM API Management provides the following key features and benefits:

- Secures, scales, and controls access to APIs to provide a resilient and flexible API runtime infrastructure. IBM API Management is powered by the DataPower Gateway appliances, which are some of the industry-leading security and integration gateway appliances.
- Empowers companies with the insight to change and grow their business in the new web API economy with robust business analytics.
- Rapidly addresses business demands with the creation of new APIs from existing business assets, or with simple configuration-based cloud services integration.
- Nurtures innovation by building a community that attracts developers, entrepreneurs, and partners who will rapidly build new applications and extend the value of the core enterprise assets.

Mobile development scenarios have a strong dependency on the quality and consumability of the incorporated APIs, and can be negatively affected when these APIs are not well-designed and consumable. Figure 5-14 depicts, at a high level, the usage of a well-defined API strategy to make the business data that is in an enterprise accessible and consumable by various consumers, such as customers, partners, vendors, and others.

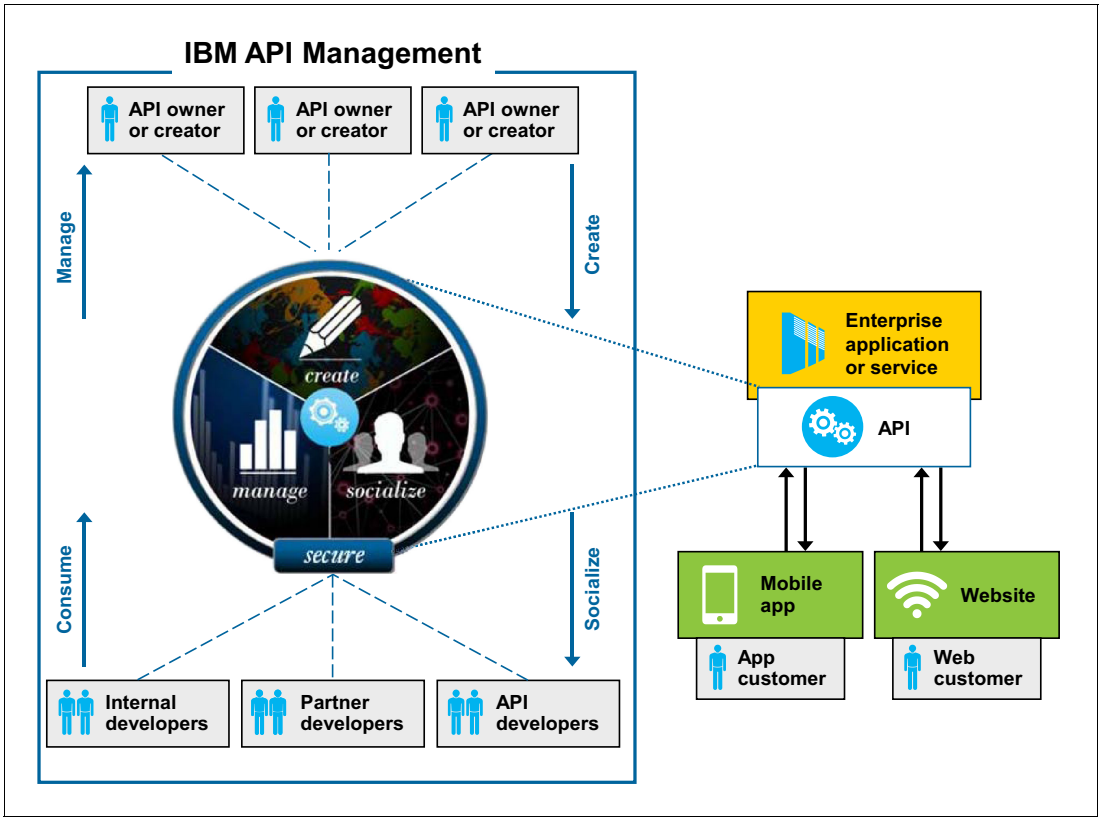

Figure 5-14 IBM API Management overview

IBM API Management provides detailed analytics and operational metrics to the business owner, and a customized developer portal for socializing the APIs and managing applications that can be used by developers. Mobile enablement of SAP business modules can use these enhanced IBM API Management capabilities, because SAP does not provide a comparable product within their portfolio.

Custom development of such a web API capability as part of a mobile project is not practical, and can easily consume the time and budget of a typically lean mobile project.

A better approach is to use the ready-to-use capabilities of an existing product, rather than developing such features for each mobile solution separately. The need for API management capabilities for a specific mobile customer scenario should be evaluated early in the project (for example, during a discovery workshop).

5.4.2 Enhancing mobile architectures by adding IBM mobile analytics and quality assurance capabilities

Scenarios for mobile technology integration implementation focus on the quality and usability of the application. Especially in the mobility domain, a negative user perception can become an exclusion criteria for the mobile application, resulting in negative community feedback and ultimately the loss of revenue.

IBM Tealeaf® is part of the IBM MobileFirst portfolio of offerings. It is an extension to the Worklight Platform, and provides ready-to-use insight into the customer's usage patterns of the mobile application. It is able to automatically detect obstacles and issues that a customer is facing when using the app. It covers the complete user interaction portfolio, from simple clicks to complex gestures.

IBM Tealeaf offers extended visibility for smartphones and tablets by capturing device-level and in-screen behaviors, such as scroll, swipe, and other actions. IBM Tealeaf can help organizations discover how users interact with their mobile features, so that they can make more informed mobile investment decisions.

IBM Tealeaf enables any mobile project running on the Worklight Platform to translate customer feedback into actionable improvements.

When implementing mobile enablement for business processes that have a large SAP footprint in the ERP landscape, it is likely that enhanced analytics requirements are placed on the mobile solution.

The SAP business audience is used to having business warehouse-like analytics for the activities within the SAP business modules. Alternatively, using traditional warehousing or analytics technics is not efficient in mobility projects, because of their complexity, and because of the different nature of systems of engagement provided by IBM MobileFirst compared to the systems of record provided by SAP.

Having ready-to-use mobile business analytics embedded in the mobile platform can be a key differentiator to enable delivery of successful mobile solutions.

Figure 5-15 shows one example of a business analytics heat map. It gives insight into which areas are used most often by consumers, and where critical situations arise.

Figure 5-15 Analyze user behaviors with visual heat maps

The heat map clearly shows the specific customer behavior. In the same way, it can analyze the usage of links and forms to better understand how the user interacts with the mobile application.

Another important tool in mobile analytics is user sentiment analysis: How consumers evaluate or rank the mobile apps. By analyzing consumer reviews, comments, and ratings, developers can get an early alert if one or more apps have problems. The Worklight Quality Assurance feature provides a powerful analytical tool to tap into the app store ranking system. Worklight Quality Assurance enables teams to capture tester and live user experience, to continuously build and deliver high-quality mobile apps.

In a fragmented and complex mobile environment, this product provides quality assurance for mobile applications, with user feedback and quality metrics available at every stage of the app development. This product also includes capabilities for validating apps and tracking production usage.

5.4.3 Enhancing mobile architectures by adding secure offline capabilities

Mobile implementation scenarios that incorporate business data are likely to have a requirement to store such business data on the mobile device itself. In some mobile scenarios, you might want a certain grade of offline capability. This might be to overcome unstable network situations, or to completely enable a business procedure to be run offline and then be consistently replicated and synchronized with the back-end system.

For such advanced requirements, Worklight provides a ready-to-use technology called *Worklight JSON store*. This component is available for the mobile applications developed on the Worklight Platform, and offers a rich set of client-side APIs to define and interact with a data store on the mobile device.

This interaction is platform-independent, and enables the use of the same code across different underlying mobile devices, which is especially beneficial when developing hybrid mobile applications.

The Worklight JSON store component also provides enhanced features, such as encryption of the stored data, and is used by multiple components of the Worklight Platform:

- The Worklight security framework uses the Worklight JSON store encryption capabilities to provide an offline authentication mechanism to the mobile application developers.
- The Worklight adapter framework uses the Worklight JSON store to replicate business data that has been called by the Worklight adapter. This approach provides a certain level of offline capability for calls into systems of record when needed in the mobile scenario.

Figure 5-16 shows how the Worklight JSON store technology on the mobile device interacts using the Worklight adapters with the ERP system in the back office.

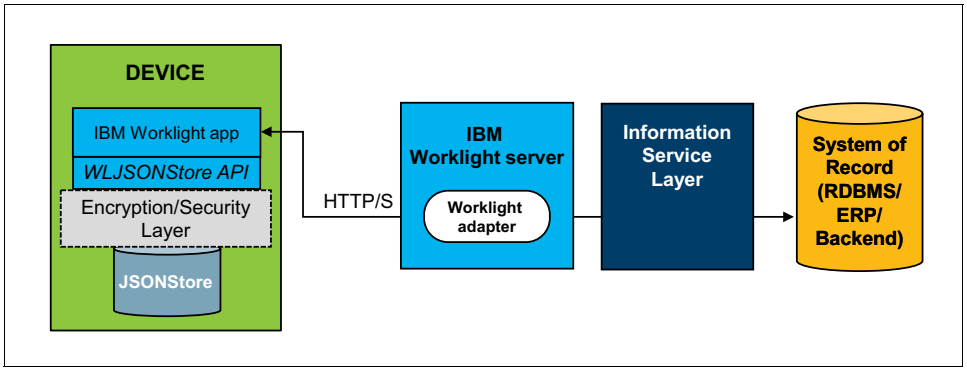

Figure 5-16 Built-in offline capability using Worklight JSON store

The Worklight offline capabilities provided by the Worklight JSON store are beneficial in SAP scenarios, because they provide a ready-to-use capability to synchronize SAP business data to the mobile device in a secure and reliable manner. The mobile developer can focus on the mobile scenario, and use the built-in framework rather than building a custom business data replication logic.

Having a product catalog or client data on hand on the mobile device while not connected to the corporate network are common requirements for mobile solutions.

5.5 Lessons learned from actual projects

This section describes some tangible findings and benefits gained from applying the IBM MobileFirst for SAP architecture to mobile development projects.

5.5.1 Direct connectivity from mobile applications to SAP is rarely used

Scenarios where the mobile application calls the SAP components directly are rare in heterogeneous enterprise architectures. Although technically feasible, such direct (unmanaged) scenarios are likely to be difficult from the operations and maintenance perspective. In addition, they provide only a limited insight into their usage patterns and customer acceptance.

5.5.2 Late decision on native versus hybrid apps

Worklight provides various development approaches, as described in the previous sections. When performing mobile application development for an SAP-driven business process, it is likely that the business data is provided by various technology stacks. In this context, Worklight adds major advantages:

- The SAP business modules offer a broad range of technologies to interact with the SAP business data, which results in different integration options on the consumer side, ranging from open standards-based integrations to SAP-proprietary approaches.
- The Worklight platform provides the mobile developer the freedom of choice to do native or hybrid development with various nuances, as is shown in 5.2, “Spectrum of mobile app development approaches” on page 119.

This flexibility is important for mobile projects where diverse SAP interface technologies are in use, and where shifts in these technologies can arise during the project phases. A change in the decision on the SAP integration interface, and changes in the mobile application design approach (for example, native versus hybrid development), can be accommodated at later stages of the mobile enablement project.

5.5.3 Adding mobile business analytics features dynamically

When delivering mobile projects for line-of-business units that are used to having SAP applications in the back-end, it is likely that business requirements arise that go into the area of mobile business analytics. These groups are used to receiving reports and evaluations, for example, *who used the business processes* and *how the customers are interacting with the mobile app*.

Adding these capabilities manually into each mobile app is a cumbersome procedure, and can increase the overall sizing for a mobile solution significantly. With the availability of the IBM Tealeaf CX Mobile capability in the IBM MobileFirst portfolio, it is fairly easy to provide such functionalities with a limited amount of custom coding.

5.5.4 Separation of security domains

SAP business modules include a highly sophisticated security model that enables business users to have various roles and permissions assigned to them. Replicating the SAP security model externally is not recommended, because the challenge to keep these models in sync can become a huge burden for any integration project.

In various projects, two major security integration patterns have been identified when interacting between mobile apps and SAP business modules.

A common security integration pattern for system-to-system integration uses a single common security context (*technical ID*) for all interactions between systems. If this pattern is applied to SAP integration and a technical ID is used to call an SAP business module, it is not obvious who in the front end is the initiator of the specific interactions. Although this approach is suitable and efficient for general system-to-system communication, it can be problematic for certain integration scenarios where a business function in SAP must be run by a user.

This problem leads to the second alternative that is called *named user* approach where the SAP transaction has to be aware of the user’s SAP ID. Many enterprises do not directly expose SAP IDs to their employees, and instead expect them to use a common enterprise-wide identity to log in to all enterprise systems, including SAP.

In this setup, the Worklight security framework, together with the mobile integration layer, determines which SAP user ID should be used to represent the mobile user and trigger the interaction in the SAP ERP system, in the specific SAP names user context. Using this approach gives the mobile solution the ability to use the full SAP security methodology. The challenge is to define, case by case, the mapping rules regarding how the mobile identity is related to the SAP back-end identity.

Establishing trust relationships between Worklight, the IBM mobile integration layer, and the SAP ERP components is the suitable solution for this common security requirement.

5.6 References

These websites are also relevant as further information sources:

- IBM MobileFirst solutions

  http://www.ibm.com/mobilefirst/us/en/offerings/

- IBM MobileFirst Platform Foundation

  http://www.ibm.com/software/products/en/mobilefirstfoundation

- Tealeaf CX Mobile

  http://www.ibm.com/software/products/en/cx-mobile

- WebSphere Cast Iron Cloud integration

  http://www.ibm.com/software/products/en/castiron-cloud-integration

- Cast Iron: Overview of the SAP connector

  http://pic.dhe.ibm.com/infocenter/wci/v6r3m0/topic/com.ibm.websphere.cast_iron.doc/SAP_Overview.html

6

Portal integration with SAP

IBM WebSphere Portal is the leading product from IBM for delivering relevant, personal, and engaging experiences. By integrating best-in-class business applications (apps) from SAP with leading digital experiences from IBM, an organization can compete more effectively and enhance the productivity of its workers.

This chapter describes the use cases, integration architectures, and guidelines for integrating SAP applications into WebSphere Portal.

This chapter includes the following topics:

- 6.1, “Overview of integrating IBM WebSphere Portal with SAP applications” on page 144
- 6.2, “Architecture overview” on page 144
- 6.3, “Types of use cases” on page 146
- 6.4, “WebSphere Portal integration with SAP app use cases” on page 147
- 6.5, “Service-level integration” on page 153
- 6.6, “Architecture guidelines” on page 156
- 6.7, “Summary” on page 157

© Copyright IBM Corp. 2014, 2015. All rights reserved.

6.1 Overview of integrating IBM WebSphere Portal with SAP applications

WebSphere Portal provides an enterprise web portal that helps companies deliver a highly personalized, social experience for their customers. Integrating business applications from SAP with WebSphere Portal, an organization can compete more effectively by providing a highly personalized, single point of access to the applications, services, information, and social connections that customers and employees need.

Several goals and constraints are involved when introducing an SAP application into an environment with an existing set of well-established user environments, applications, and processes. In some cases, these items can be in conflict with one another, which requires you to strike a balance to ensure that the business objectives are met without compromising the integrity and efficiency of the architecture.

Different types of SAP application users require different levels of integration. Internal SAP application power users use the native SAP application graphical user interface (GUI) or intranet portal. Casual internal SAP application users access functions through the intranet portal, which wraps content from the SAP NetWeaver Portal component (users do not access the SAP NetWeaver Portal component directly).

Customers and business partners use an external portal that wraps the SAP application functions but provides company-standard branding, with a non-SAP application UI look and feel. Mobile applications are focused on a specific business function, for example, labor claiming.

6.2 Architecture overview

Figure 6-1 on page 145 provides an overview of how WebSphere Portal can effectively integrate SAP applications, such as SAP customer relationship management (CRM), supplier relationship management (SRM), supply chain management (SCM), Product Lifecycle Management (PLM), and enterprise resource planning (ERP), with SAP NetWeaver. It also integrates with non-SAP enterprise and cloud applications.

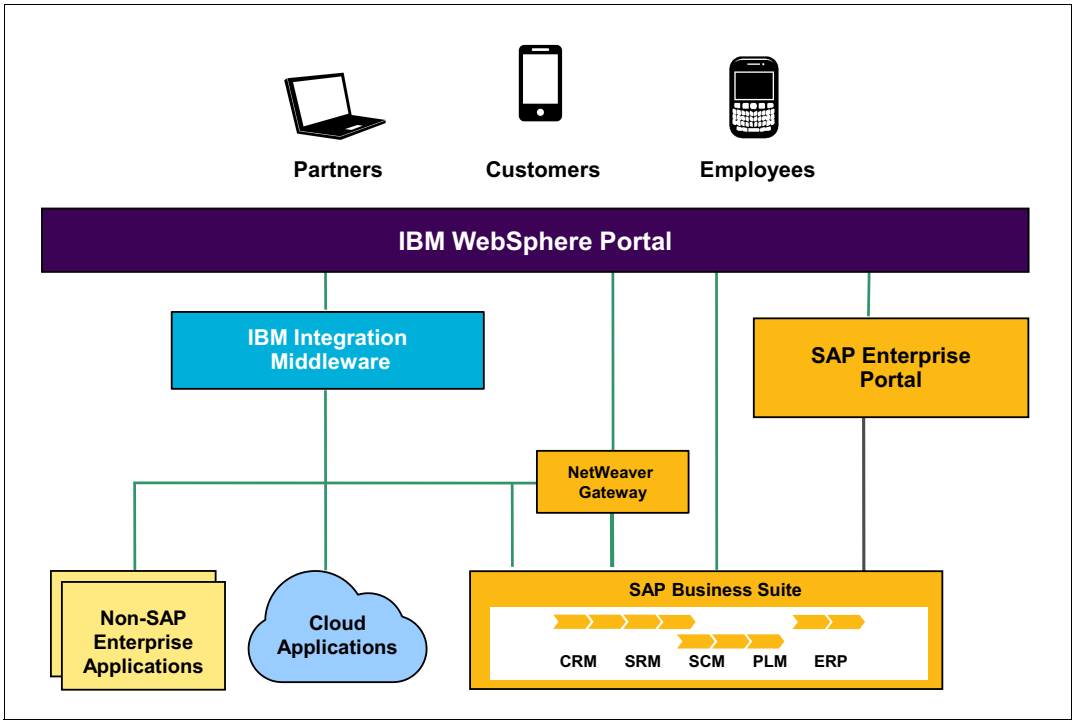

Figure 6-1 Architecture Overview diagram

WebSphere Portal helps organizations create highly engaging, personalized, and differentiated digital experiences that meet the evolving needs of customers and employees. It can help organizations deliver exceptional digital experiences:

- Interact with the appropriate back-end applications, such as SAP applications, to extend core business processes to all users and customers.
- Provide relevant, highly personalized experience according to the user's preferences, behaviors, location, and device.
- Deliver consistent experiences across multiple online channels.
- Engage through online communities, social interaction, and collaboration.
- Empower business owners to manage the creation and delivery of rich content.
- Deliver engaging experiences without sacrificing flexibility, scalability, or security.

6.3 Types of use cases

SAP application use cases can be categorized into *casual* and *detailed*, as shown in Figure 6-2.

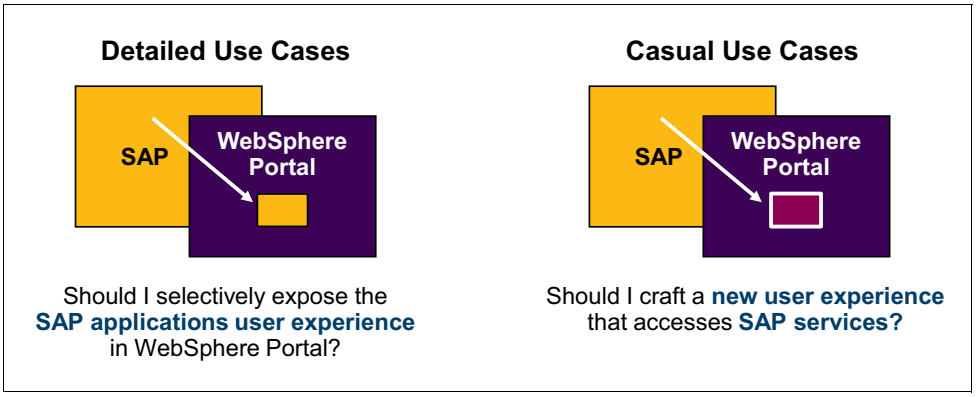

Figure 6-2 Types of SAP applications use cases

6.3.1 Casual use cases

The majority of employees and possibly customers will encounter casual use cases. These users need occasional access to services and information that originates from SAP applications. They need the information in the context of what they are doing, and do not need to know or care that an SAP application is involved. An example of a casual use case is a salesperson looking up customer information or pricing. Another example can be a customer who has been provided with visibility into their invoices and billing.

In both cases, the requirement is to provide simplified access to SAP application content in the context of their role and task at hand. Casual use cases are often best addressed by a new or simplified component that integrates with the SAP application at a service level, providing the specific information needed in the context and manner that is most appropriate for the user.

It also bypasses any complex SAP application screens that might detract from the usability of the experience. The casual use case is common in enterprises that have many user interfaces across heterogeneous environments.

6.3.2 Detailed use cases

Detailed use cases typically involve more than simple access to content:

- A sales person creating a new customer opportunity in the CRM system
- A pricing analyst managing product and pricing data
- A manager performing salary planning for his team

The SAP application provides a ready-to-use user experience that has been refined to meet the needs of each of these scenarios. The experience ties directly into logic on the back end, and involves some level of interaction with the user.

These experiences are delivered by SAP software as *business packages* for each of the various SAP business applications. Business packages are sets of the SAP NetWeaver Portal component iViews (visualized in a web browser), combined with back-end enhancements that *snap* into the SAP NetWeaver Portal component.

For detailed use cases, the approach that makes most sense is to reuse, and use the experience that the SAP NetWeaver Portal component provides by exposing it in WebSphere Portal (in addition to user experiences from other systems). The SAP NetWeaver Portal component experience should be exposed in a way that makes it feel like a natural part of WebSphere Portal.

The experience must be integrated in the context and user's role, and not require a separate SAP NetWeaver Portal component sign-on. Ideally, the SAP NetWeaver Portal component experience is exposed in WebSphere Portal in a way that it is transparent to the users which systems they are working on.

6.4 WebSphere Portal integration with SAP app use cases

One size does *not* fit all regarding SAP application integration. Different users and use cases are best served by different types of integration. The types of integration can be divided into the following categories:

- Expose and reuse SAP user experience inside of WebSphere Portal.
- Create a new user experience, or a new use of a non-SAP experience, that accesses SAP Services.

The following sections describe common integration use cases.

6.4.1 Federated portal

This type of integration is ideal for detailed use cases. The SAP NetWeaver Portal component user experience is reused by exposing it in WebSphere Portal. Various features are available in WebSphere Portal to achieve this goal, and the decision to choose the most appropriate feature depends upon whether you want to selectively expose specific content pieces from the SAP NetWeaver Portal component, or you want to expose the complete SAP NetWeaver Portal component content.

Tip: For branding alignment between SAP and IBM Portal products, organizations should build equivalent themes in the SAP NetWeaver Portal component using the Theme Editor tool, and in WebSphere Portal using the Theme Builder tool. Aligning the themes between the portals is essential to provide users a seamless and harmonized integration.

6.4.2 Integrating with the web application bridge feature

The web application bridge (WAB) is a feature of WebSphere Portal that uses reverse proxy technology to integrate different web applications natively inside WebSphere Portal pages so that WebSphere Portal becomes the single window of access to users. Organizations can use WAB to quickly integrate the SAP NetWeaver Portal component into WebSphere Portal to enhance their user experience.

WAB provides *at the glass* integration, and can be used to integrate web applications based on multiple technologies and multiple vendors. It is installed and ready to use in WebSphere Portal version 8.0.0.0 and later. It is also available for WebSphere Portal version 7.0.0.2 in the form of a Portal Application Archive (PAA) file, which can be downloaded at no charge from the IBM Collaboration Solutions Catalog at the following website:

`http://ibm.co/1mLOMcc`

Organizations can use this solution today, without having to purchase any additional software. WAB is based on HTML iFrames and reverse proxy technology, and consists of the following components:

- Virtual Web Application Manager portlet. An administration portlet that provides a centralized management console for all of the applications that are being integrated through WAB.
- Web Dock portlet. An advanced iFrame-based portlet that can dynamically resize content (no scroll bars). It enables client-side or server-side inter-portlet communication, session alignment, and navigation state saving.
- Engine. A back-end component to manage persistence of the WAB configuration.
- Reverse proxy servlet (RPS). The servlet that proxies every Hypertext Transfer Protocol (HTTP) request that is served through the Web Dock portlet's iFrame, and processes single sign-on (SSO).

Figure 6-3 depicts the general flow of an HTTP request that is served when a web application is integrated using WAB.

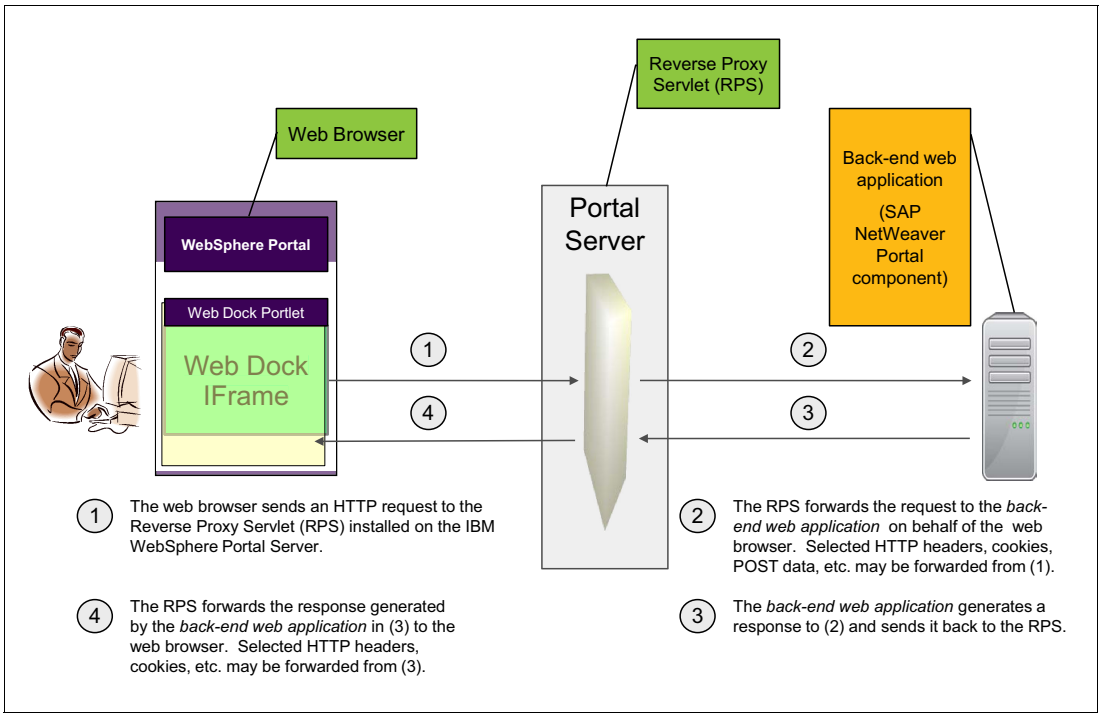

Figure 6-3 Request/response flow in web application bridge

WAB provides the SSO capability to enable the users to access the integrated SAP NetWeaver Portal component content by logging in just once into WebSphere Portal. *Basic* and Security Assertion Markup Language (*SAML*)-based authentication are the two common authentication mechanisms used by the SAP NetWeaver Portal component, and both of them are supported SSO mechanisms in WAB.

Note that for SAML support in WAB, the SAML token should be inside a cookie, and the portal server and the target server should be in the same domain.

More information: For details regarding using WAB to integrate web applications into WebSphere Portal, see the *Integrating with web applications* web page:

`http://www.ibm.com/support/knowledgecenter/SSHRKX_8.5.0/mp/admin-system/wab.dita`

With the SAP interoperability framework page, shown in Figure 6-4, you can render the SAP NetWeaver Portal component content without the navigational elements, so that it is cleanly exposed in WebSphere Portal.

More information: For information about how to construct a Uniform Resource Locator (URL) for an iView or page when using the interoperability framework, and then configure that constructed URL using the interoperability framework into the Web Dock portlet, see the "The SAP Interop Framework Page" section in the *Defining Consumption Mode and Creating the Content Component* web document:

`http://help.sap.com/saphelp_nw73/helpdata/en/f5/9edaa160584fb59081fef067b7a415/content.htm`

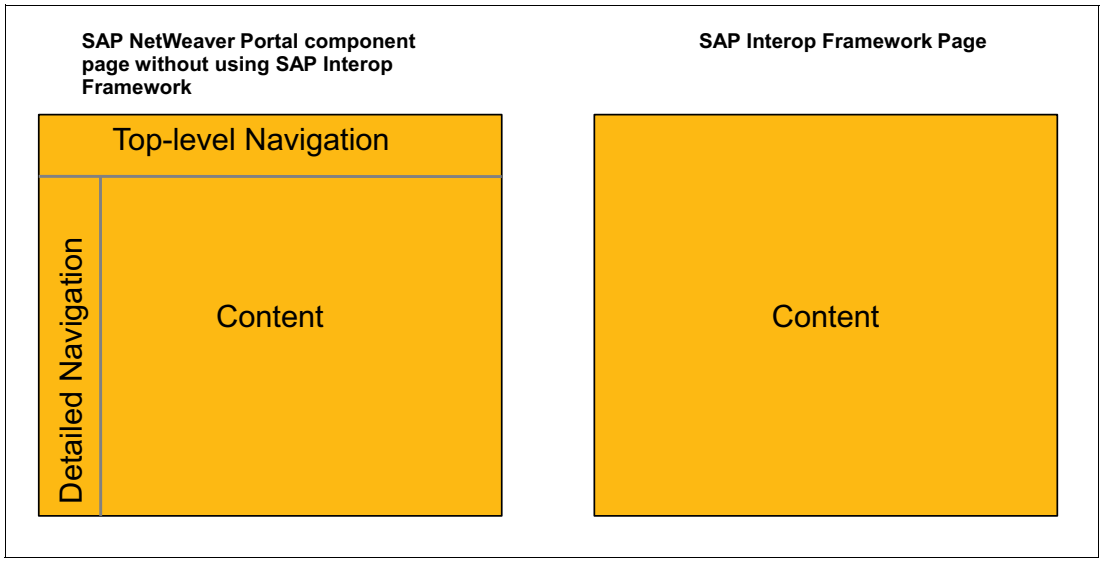

Figure 6-4 SAP Interop Framework Page

A key challenge in integrating the SAP NetWeaver Portal component is that when a user logs out of WebSphere Portal, a log out of the SAP NetWeaver Portal component needs to be performed also. Otherwise, the user session on the SAP NetWeaver Portal component remains open until it times out. WAB includes a JavaScript-based plug-in to handle this situation, so that writing any custom code or performing any configuration is unnecessary.

The SAP NetWeaver Portal component sessions that are started in the SAP NetWeaver Portal component and related SAP back-end applications are aligned with WebSphere Portal. This way, the required session control and management are possible.

In most cases, session state is maintained if users go to another section of the portal. Therefore, when they return to the SAP NetWeaver Portal component content, they can resume access to the SAP NetWeaver Portal component. When a user logs out of WebSphere Portal, the related SAP NetWeaver Portal component session will be closed. In all cases, the session is closed when the user closes the browser.

Using WAB to integrate the SAP NetWeaver Portal component, SAP application content can be placed alongside the information from other systems, including web content and social capabilities from IBM. The WAB also provides support for JavaScript-based, client-side, inter-portlet communication, in addition to the standard Java Specification Request (JSR) 286-based server-side eventing.

More information: For details about WAB inter-portlet communication, see the “Web application bridge inter-portlet communication” topic at the following location:

`http://www.ibm.com/support/knowledgecenter/SSHRKX_8.5.0/help/panel_help/h_wab_ipc.dita?lang=en`

6.4.3 Integrating with IBM WebSphere Portal Integrator for SAP

IBM and SAP collaborated to provide integration capabilities that enable the SAP NetWeaver Portal component to interoperate with WebSphere Portal. This integration capability is well-suited for the use cases in which an organization wants to expose an entire set of the SAP NetWeaver Portal component content, such as all of the pages and their associated navigation.

The IBM WebSphere Portal Integrator for SAP is a feature of WebSphere Portal available, starting from version 7.0.0.1 Cumulative Fix 6 (CF6) onwards, that provides interoperability with the SAP NetWeaver Portal component. It is based on new public SAP application programming interfaces (APIs) and new features introduced in SAP NetWeaver Portal 7.3, and is jointly supported by IBM and SAP.

This new feature is available at no charge to all WebSphere Portal customers, and is available in the IBM Business Solutions Catalog. SAP NetWeaver Portal 7.3 (Enterprise Portal Core minimum) customers and WebSphere Portal (version 7.0.0.1 CF6 and later) customers can implement this solution today without having to purchase any additional software.

The portal interoperability solution addresses several main technical requirements for seamless portal interoperability:

- SSO
- Navigation federation
- Session management

SSO between WebSphere Portal and the SAP NetWeaver Portal component are handled by WebSphere Portal Integrator for SAP through either basic authentication (using the credential vault) or with SAML 2.0. Using the credential vault approach is quick and easy, but also has a downside, because it requires the user to provide the user ID and password redundantly to WebSphere Portal and the SAP NetWeaver Portal component. The credential vault is typically used in proofs of concepts (PoCs) and technical evaluation scenarios.

Production-level security integration with automatic SSO, without any additional need for the user to provide data, can be provided by SAML or similar SSO infrastructures solutions, such as IBM Security Access Manager. These solutions typically include some form of user ID mapping as required by the two systems’ infrastructure.

Consumption of the SAP NetWeaver Portal component navigation structure is achieved with a new public SAP NetWeaver Portal component navigation web service, which exposes the user navigation structure through a web service.

As shown in Figure 6-5, the new WebSphere Portal Integrator consumes the navigation structure from the logged-in user of the SAP NetWeaver Portal component, and seamlessly integrates into that user's session in WebSphere Portal. The result is a federated navigation structure for the user that takes into account the SAP NetWeaver Portal component content that has been integrated.

WebSphere Portal Integrator for SAP achieves seamless integration with the SAP NetWeaver Portal component user experiences in WebSphere Portal by performing the following tasks (Figure 6-5):

- Providing SSO from WebSphere Portal to the SAP NetWeaver Portal component. The SSO flow is performed in the background, and users are not aware that they are logged in to the SAP NetWeaver Portal component.
- Consuming the SAP NetWeaver Portal component navigation structure for the user and role into WebSphere Portal, and displaying the associated content.

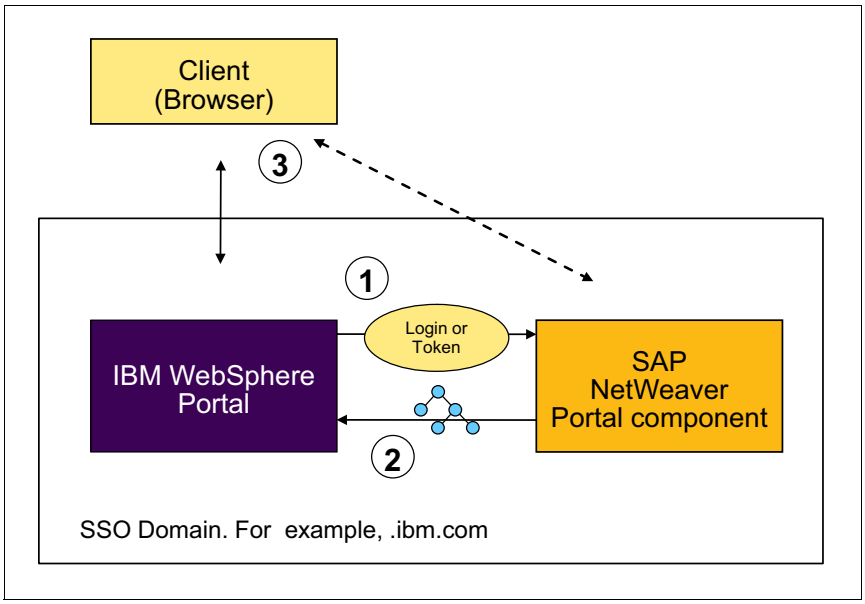

Figure 6-5 Integration with WebSphere Portal Integrator for SAP

The user logs in to WebSphere Portal and appears to be working in a single integrated application. In reality, the user is actually logged in to two different systems, interacting directly with the SAP NetWeaver Portal component and working in the SAP NetWeaver Portal component content. With this approach, everything behaves exactly as it should in the SAP NetWeaver Portal component.

All the servers must be part of the same SSO domain; otherwise, the cookies will not be handled correctly by browsers. An important step is that the users specify the full SSO domain when accessing the systems. Users must have direct access to all of the involved servers. Of course, a proxy can be used.

The SAP NetWeaver Portal component sessions that are started in the SAP NetWeaver Portal component and related SAP back-end applications can be aligned with WebSphere Portal. In this way, the required session control and management are possible.

In most cases, session state is maintained if users go to another section of the portal. Therefore, when they come back to the SAP NetWeaver Portal component content, they can resume access to the SAP NetWeaver Portal component.

When a user logs out of WebSphere Portal, the related SAP NetWeaver Portal component session is closed. In all cases, the SAP NetWeaver Portal component will end the session when the user closes the browser.

WebSphere Portal Integrator for SAP can integrate the SAP NetWeaver Portal component content in a way that makes it feel like a natural part of the WebSphere Portal user experience. Content can be placed alongside information from other systems, including web content and social capabilities from IBM. This capability enables reuse of the SAP NetWeaver Portal component investment by exposing it in the social business context of WebSphere Portal, thereby providing maximum value and return on investment (ROI).

6.4.4 Integrating with Web Services for Remote Portlets (WSRP)

WebSphere Portal supports the Web Services for Remote Portlets (WSRP) standard. By using this standard, a portal, such as the SAP NetWeaver Portal component, can provide portlets, applications, and content as WSRP services. Other portals, such as WebSphere Portal, can then integrate the WSRP services as remote portlets for its users.

The WSRP standard and specification is provided by the Organization for the Advancement of Structured Information Standards (OASIS). It defines a web service communication interface for interactive presentation-oriented web services. This standard simplifies the integration of remote portlets, applications, and content into portals. Producers and consumers use this interface for providing and consuming portlets. WSRP can be used in the following ways:

- Producers can provide portlets as presentation-oriented WSRP services, and make them available to consumers that want to use these services.
- Consumers can select from a rich choice of available remote portlets, and integrate them into their portal.
- Portal site visitors can then access the integrated remote portlets. They can work and interact with them in the same way as they do with local portlets. The integrated remote portlets appear and operate to portal site visitors the same as local portlets.

More information: For more information about the OASIS WSRP specification, see the *OASIS Web Services for Remote Portlets (WSRP) TC* web page:

`https://www.oasis-open.org/committees/tc_home.php?wg_abbrev=wsrp#overview`

Note that the SAP NetWeaver Portal component provides limited support only for the case where it is being used as a WSRP producer. Specifically, the SAP NetWeaver Portal component WSRP production does not support the following aspects of WSRP:

- SAP NetWeaver Portal component role, workset, and page concept
- SAP NetWeaver Portal component business packages
- Web Dynpro
- iViews started by the SAP Application Integrator, which computes a URL to the application and uses browser redirect to start it
- Application-oriented services:
  - Client eventing
  - Object-Based Navigation (OBN)
  - Work Protect mode to secure the user data

Effectively, WSRP-based integration of pre-built contents of the SAP NetWeaver Portal component is not supported, but one can develop custom Java-based iViews and make them available for WSRP consumption.

The WSRP implementation of WebSphere Portal version 8.5 is built on the Java API for Extensible Markup Language (XML) Web Services (JAX-WS)-based web service stack of IBM WebSphere Application Server. The WSRP implementation in earlier versions of WebSphere Portal was based on the web service stack that is based on the JAX-based Remote Procedure Calls (JAX-RPC) standard. WSRP in WebSphere Portal V 8.5 can interoperate with WSRP counterparts that are built on other web service stacks, such as JAX-RPC.

6.5 Service-level integration

This type of integration is ideal for casual use cases. User experiences can be easily crafted to access SAP services, and provide the information in the context and level that are most helpful. IBM provides options for creating user experience components that access SAP services. The primary tool for this purpose is IBM Web Experience Factory, which accelerates the delivery of enterprise-ready, standards-based, Web 2.0 applications with rich, interactive interfaces that deliver exceptional digital experiences.

Developers of all skill levels can use IBM Web Experience Factory to rapidly create complex, rich, and interactive applications by using the latest Web 2.0 technologies, including Ajax and advanced Dojo toolkit widgets and controls. IBM Web Experience Factory provides flexible deployment options with support for the most popular smartphone devices, including Apple iPhone, Android, and BlackBerry, all from a single code base.

At the core of IBM WebSphere Experience Factory Designer is a set of software automation components called *builders*. These builders capture design intelligence through easy-to-use, wizard-like user interfaces, and then automate the creation of code. The use of builders greatly speeds up the development process, thereby masking the complexities of the underlying APIs, and producing portlets that are truly compliant with service-oriented architecture (SOA).

Included in the IBM Web Experience Factory are several builders that enable developers to rapidly and easily integrate with SAP applications to create applications based on any remote-enabled SAP Business Application Programming Interface (BAPI), Remote Function Call (RFC), web service, or Representational State Transfer (REST) call. By using these builders, companies can speed up time-to-market for new SAP application user experiences to address casual use case requirements.

IBM Web Experience Factory has the following key features:

- Fully automated SAP applications integration. Easily created, enterprise-quality, custom portlets and applications that enable users to search, view, sort, create, update, and delete SAP data.
- Robust personalization and customization capabilities. Enable business users and end-users to personalize and customize any aspect of their SAP applications, including the look and feel, data, and application flow.
- Dynamic profiling. Enable adaptive applications that display different data, enable different tasks, and enable different administrative rights depending on the role or group of the user (for example, administrator, employee, partner, and so on).
- Automatic translation of help values. Effortlessly translate SAP application codes (for example, country, currency, and so on) into user-friendly text, select lists, or radio buttons.

- Comprehensive Object Browser. Quickly explore all BAPIs within the SAP Business Object Repository. Drill down to view the object details, such as methods, parameters, structure, and field attributes.
- Globalization. Easily build globalized portlets by using IBM Web Experience Factory's ready-to-use support for resource bundles, multi-byte characters, and runtime selection of locale-specific content.
- Simplified portlet-to-portlet communication. Create a richly integrated portal experience by enabling portlets to interact, even if they are accessing data from disparate databases and systems.
- Batch input support. Rapidly import data into SAP applications using recorded transactions.

6.5.1 Direct integration with SAP applications using SAP Java connector

This section describes the direct integration with SAP applications using SAP Java connector (SAP JCo) to access SAP RFC-enabled BAPIs. This is a quick and easy approach to create portlets that work with SAP applications. Anyone with access to an SAP server can browse and directly access the SAP RFC-enabled BAPIs. New user experiences can be rapidly created and deployed to meet changing business requirements.

Web Experience Factory builders work with =SAP JCo to access RFC-enabled BAPIs. SAP JCo is a component provided by SAP for Java-based development of SAP-compatible applications. The Web Experience Factory builders provide full support to create, read, update, and delete (CRUD) SAP information. The SAP RFC-enable integration model acts as a service provider to one or more user experience models.

This approach offers flexibility because you can reuse the data as you refine the user experience and build new ones. By separating SAP application integration from the user experience, it also buffers the user experience from any changes on the back-end SAP system (see Figure 6-6).

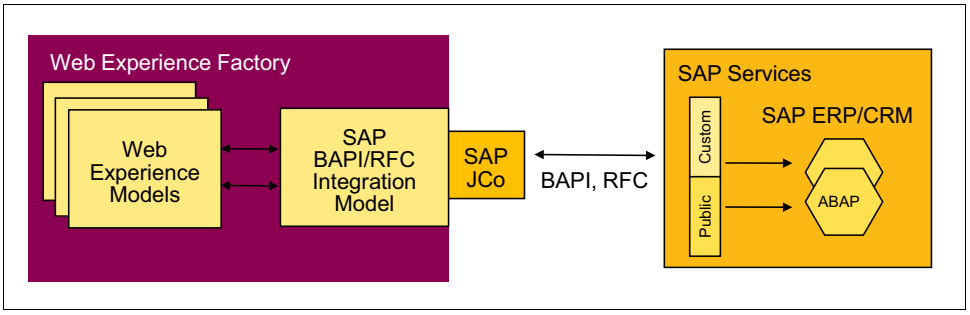

Figure 6-6 Integration using SAP JCo through RFC-enabled BAPIs in SAP

The SAP view and form builder work with the BAPI builders to provide ready-made input and output experiences that can serve as a basis for further customization.

User credentials can be passed to SAP applications through the IBM Web Experience Factory builders, thereby enabling you to create a custom experience that accesses SAP without the user even knowing it. This token-passing infrastructure can use SSO infrastructure solutions as provided by SAML, or other token-passing SSO solutions, such as IBM Security Access Manager. For simple PoC scenarios, it can be used with the credential vault SSO solution provided by WebSphere Portal.

6.5.2 Integrating with an enterprise service bus to connect to SAP applications

A second approach to service-level integration is to use an enterprise service bus (ESB) to connect to SAP applications. This approach is common where an SOA architecture is being pursued. The service bus acts as a middleware buffer, which provides access to SAP and non-SAP systems through web services.

It simplifies the integration of SAP applications by providing a single connection that is shared and managed through the ESB. This approach requires more infrastructure work to set up, but it can provide significant SOA benefits in the long term.

The IBM Web Experience Factory architecture includes a web service model in place of a direct connection through RFC-enabled BAPIs. The user experience model remains the same. Changes to the back-end are buffered through this web services layer, providing flexibility to release level changes and updates to SAP applications. The user ID mapping and SSO is typically provided by the ESB in this approach, as shown in Figure 6-7.

Any enterprise with many heterogeneous systems driving toward existing modernization, SOA, business process management (BPM), and mobile technologies should seriously consider this architecture.

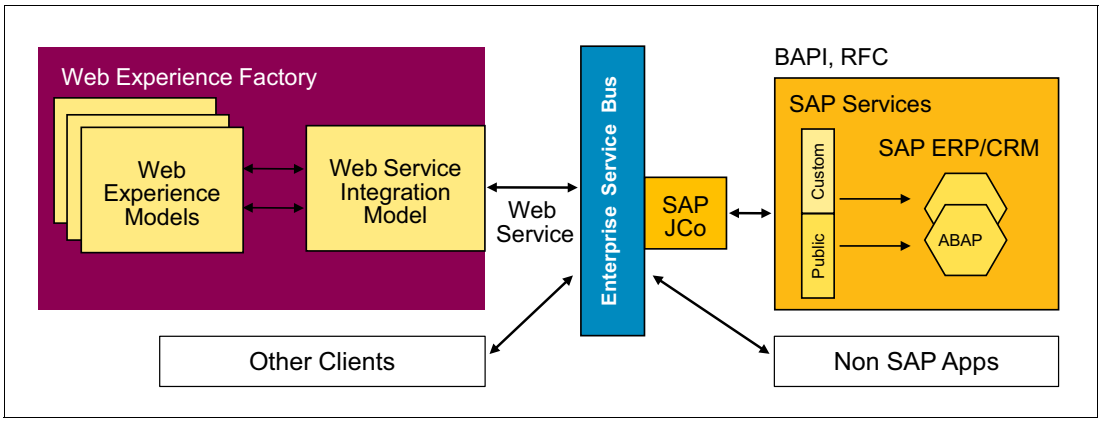

Figure 6-7 Integration using an ESB

6.5.3 Integrating with SAP NetWeaver Gateway

A third approach to service-level integration is through IBM Web Experience Factory integration with SAP NetWeaver Gateway (that is, REST and OData services from SAP). This approach offers connectivity to SAP applications, without the need for specific SAP knowledge, by using REST services.

IBM Web Experience Factory developers can build REST SAP application integration models without having detailed knowledge of BAPIs and RFCs. Although this pattern does not require SAP skills to accomplish portal integration, additional and potentially complex SAP application configuration might be required in SAP NetWeaver Gateway to service-enable access to SAP application functions and data.

REST is an inherently lightweight and intuitive environment that enables developers to create, update, query, and manage information of any REST-enabled applications (such as SAP) from their own custom applications. The simplicity in the API comes from the fact that the REST API is HTTP-based, enabling you to easily make requests to SAP through a simple and straightforward URL using **`POST`**, **`PUT`**, and **`DELETE`** methods.

This new approach provides a simple and fast way to create user experiences to address casual use cases. It uses the current REST builders in IBM Web Experience Factory with SAP's strategic focus, enabling easier access to SAP from third-party products and devices, as shown in Figure 6-8.

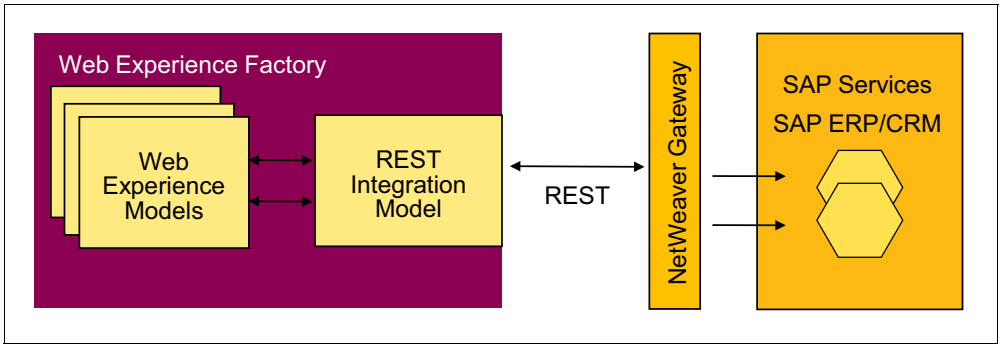

Figure 6-8 Integration using SAP NetWeaver Gateway

6.6 Architecture guidelines

The ROI on SAP applications implementations is directly linked to the level of SAP applications customization. The cost of maintaining SAP applications has a significant negative effect on the ROI.

Use of the pre-built, ready-to-use SAP application user experience is a good approach when the level of customization required to meet actual business requirements is, roughly, less than 10%. In this case, selective exposure of such pre-built SAP NetWeaver Portal component user experience inside WebSphere Portal (integration on the glass) should be done with IBM web application bridge, as described in 6.4.2, "Integrating with the web application bridge feature" on page 147.

Complete exposure of the SAP NetWeaver Portal component user experience should be done with WebSphere Portal Integrator for SAP, as described in 6.4.3, "Integrating with IBM WebSphere Portal Integrator for SAP" on page 150.

An important note is that the SAP NetWeaver Portal component should be used only as a *black box* (only its externally visible behavior is considered) runtime environment. This black box runtime environment should only host prepackaged SAP NetWeaver Portal component user experience content when it matches business needs with no or bare minimum customization.

The SAP NetWeaver Portal component should not be used when significant levels of customization are required, because it results in significant negative effect on the ROI of SAP application implementation.

Service-level integration that provides new user experience, as described in 6.5, "Service-level integration" on page 153, should be used when a specific function is not available in the prepackaged SAP application user experience contents, or when such a function requires a significant level of customization (for example, extra simplicity required for the casual business user).

Service-level integration is also the preferred option when an SAP application is only one of several back-end sources providing a similar type of information. For example, the parts search application might use the SAP application as one of the ways to locate the needed part, along with other systems that cover different areas or partners. In this case, the SAP application and other source systems should be represented as a set of services, and the application UI should be built with the service-level integration.

Selection of the specific sub-options for service-level integration (described in sections 6.5.1, “Direct integration with SAP applications using SAP Java connector” on page 154 through 6.5.3, “Integrating with SAP NetWeaver Gateway” on page 155) depends on the project context as described in the following considerations:

- Direct integration with SAP applications is best suited for small projects with isolated development of a portal web UI front end for the SAP application. This option enables developers to directly explore, select, and combine in the UI module (a portlet) a set of remote function services directly exposed by the SAP back-end application.

  This approach enables a degree of agility for fast UI implementations, because both UI and integration development falls into a single team. However, it effectively bypasses services reuse and governance. It also lacks an effective solution for user identity mapping (Portal ID to SAP ID).
- Integration through an ESB is a preferred approach for all medium-to-large SAP application integration projects.
- Integration using SAP NetWeaver Gateway uses the flexibility and interoperability of the IBM platform. It should only be considered as an alternative for ESB.

6.7 Summary

This chapter demonstrates how an organization can use existing investments in WebSphere Portal and SAP applications by integrating SAP applications into WebSphere Portal. This integration can be accomplished by surfacing the UI of the SAP application directly into the portal using UI-level integration tools and techniques.

Another option for integration is to consume various services exposed by the SAP application, and create a custom portlet-based UI in WebSphere Portal using deeper, service-level integration tools and techniques. This chapter also described how to decide which type of integration and which tools and techniques to use for various scenarios.

6.8 References

These websites are also relevant as further information sources:

- IBM WebSphere Portal family

  http://www.ibm.com/software/products/en/websphere-portal-family
- SAP Enterprise Portal (formerly known as SAP NetWeaver Portal)

  http://scn.sap.com/community/enterprise-portal
- Interoperability of SAP NetWeaver Portal 7.3 and IBM WebSphere Portal

  http://scn.sap.com/docs/DOC-26539

7

Master data management for SAP

This chapter introduces the concept of master data management (MDM) and provides an overview of the IBM Master Data Management solutions. This chapter shows why SAP applications need MDM, and why IBM MDM is a particularly well-suited MDM solution for a heterogeneous application landscape with SAP and non-SAP applications. This chapter describes the solution architecture and architecture patterns that can be applied to enable better business outcomes by extending SAP applications with IBM MDM.

This chapter includes the following topics:

- 7.1, “Master data management introduction” on page 160
- 7.2, “Why master data management is important for SAP applications” on page 161
- 7.3, “Overview of IBM Master Data Management capabilities” on page 163
- 7.4, “Architecture goals” on page 166
- 7.5, “Architecture overview” on page 166
- 7.6, “IBM InfoSphere MDM for SAP applications” on page 168
- 7.7, “Architecture patterns” on page 170
- 7.8, “References” on page 187

© Copyright IBM Corp. 2014, 2015. All rights reserved.

7.1 Master data management introduction

Master data, such as customer, supplier, employee, material, product, asset, contract, or agreement, encompasses the key entities for an organization. Master data appears in almost all relevant business processes. Master data has, therefore, high value information for an enterprise. In some enterprises, master data is not properly managed today.

Often, it exists with poor data quality, without proper management functions in application and system islands. This leads to inconsistent, duplicated master data information across the IT landscape.

MDM solves the problem of multiple and inconsistent source systems for master data by centrally managing these business-critical information assets, consistently and with the highest degree of data quality. Also, MDM has to provide the trusted master data to all relevant business processes in a timely way, either with batch-oriented or (near-) real-time business services.

Figure 7-1 shows an example of ready-to-use IBM InfoSphere Master Data Management user interface (UI) for managing customer master data.

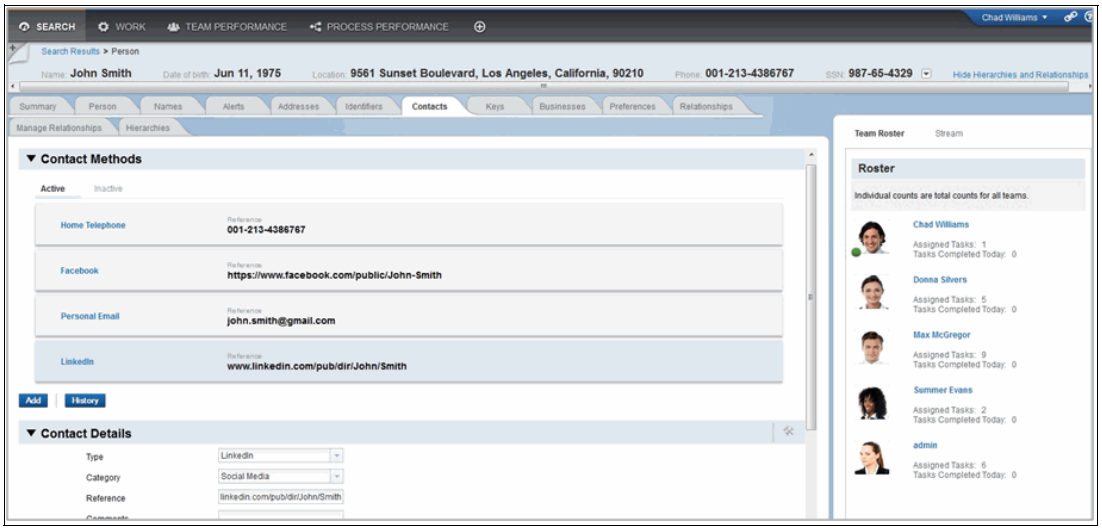

Figure 7-1 IBM InfoSphere Master Data Management: User interface example

At a high level, several major benefits are derived from the adoption of an MDM solution:

- Lower operational costs. Streamlining master data processes with an enterprise-wide MDM solution, removing redundant master data silos, and providing simplified, less-error prone master data quality management through automation reduces operational costs.
- Improved agility. Getting products to markets faster, onboarding customers faster, and being able to identify new sales opportunities more quickly make an organization much more nimble and agile to materialize better business outcomes through MDM.
- Improved risk and compliance management. By providing 360° consistent customer profiles, MDM can help to reduce fraud and improve compliance. Process-driven MDM master data is more accurate, because the processes consistently consume the same MDM services and rules governing master data.

  In addition, critical steps in the maintenance of master data cannot be skipped, because of process repeatability. This approach enables much more consistent adherence to regulations or industry standards and simplifies audit reporting.

With centralized master data and well-defined distribution to consuming applications, organizations also have a deeper understanding of where each piece of master data resides, which is a prerequisite for controlling access to it. Consistent assignment of customer preferences avoids incidentally sharing customer data with others, against the customer intent, or inappropriately contacting a customer. Centrally managed status on credit scores, black list status, and so on, mitigates business and compliance risks.

- Increased sales. With a 360° customer profile, companies can shift from an organization that is account-centric to one that is customer-centric by gaining an understanding of sales opportunities across line of business (LOB) boundaries. Also, for individual customers, shifting from an account-centric view per LOB to a customer-centric view across LOBs enables a superior customer experience across all channels by treating customers individually, in a personalized manner, and improving customer relationships.
- Social MDM. At the intersection of big data, mobile technologies, and master data, social MDM is on the rise. A complete customer profile not only covers the enterprise-internal knowledge about the customer, it also includes insight into all interactions with the customer across all channels. For example, was the customer happy last time a call center interaction occurred?

  In addition, such a complete profile includes insight into the particular features a customer liked or disliked, who is influential on key social media platforms, and so on. For example, a prospect interacts through the Facebook presence of a company to get information about a certain product. On the next day, the customer stops by the local branch office to look at the product. The potential customer expects that the sales clerk is aware that the Facebook conversation occurred, and has information about its content.

  Consequently, unstructured information needs to be processed on Apache Hadoop platforms, and derived social personas must be matched against the customer records in MDM. If matches are found, the customer profile in MDM must be enriched accordingly, enabling a complete social experience for the customer.

7.2 Why master data management is important for SAP applications

The SAP application suite provides functionality to manage the master data that SAP needs for the particular business scope, such as enterprise resource planning (ERP) or customer relationship management (CRM). However, a true MDM solution accomplishes much more than just managing a local repository of master data.

A true MDM solution is configured to be aware of all applications across the enterprise that contain master data, and to provide a layer of data integrity, referential integrity, governance, and management across all master data repositories in a cohesive manner. In addition, an MDM solution usually maintains its own local repository of master data, which has been consolidated and cleansed from all of the other master data repositories from across the enterprise.

The MDM consolidated repository does not contain every attribute from every application, but rather just the most important attributes that are common across all applications. If the MDM solution needs access to an attribute that is specific to the functional scope of one of the applications, it maintains reference links to the source application to retrieve that data as needed.

Having an MDM solution in place to complement SAP adds value to SAP by reducing the need to customize packaged SAP functionality. Without MDM, clients are often tempted to extend master data definitions within SAP applications to suit various needs of a particular business scenario. With an MDM solution in place, the MDM platform is designed to incorporate constantly evolving master data definitions, therefore avoiding the need to customize the SAP application beyond the intended scope of the solution.

For example, the master data definition for *customer* in SAP ERP Central Component (SAP ECC, an ERP solution) should not contain any more attributes than those needed for the scope of ERP (financials, order fulfillment, and so on). The master data definition for *customer* in SAP CRM should not contain any more attributes than those needed to support the sales and support processes.

Note that SAP CRM and SAP ERP have different data models for customer and product. SAP CRM, for example, has the business partner entity for managing customer data with a nice distinction between *individual*, or business-to-consumer (B2C), and *organization*, or business-to-business (B2B), customer types, based on the BUTXXX[1] table family. However, the customer data model in SAP ERP based on the KNXX[2] table family lacks this clear distinction.

Similarly, the data model for product is quite different across these two applications. These facts illustrate the need for an application-independent, enterprise-wide master data model for each master data entity.

In addition to the data model differences, SAP applications are often customized during deployment, and many enterprises have more than one SAP CRM or more than one SAP ERP, or both, deployed in regional or LOB roll-outs. Customization in SAP applications enables, for example, different reference values in code tables (also known as *check tables*), and they contain, for example, country codes, product codes, and so on.

Therefore, there is no consistent definition of master data, even within the same application across multiple instances of the same SAP application. As a net result, downstream applications, such as data warehouses, face challenges that could occur in report quality, for example, report by country, by account group, and so on.

Another angle of customization frequently applied is the addition of tables, or custom attributes on existing tables, also modifying the data model. Therefore, an MDM system external to the SAP applications that is able to support all SAP and non-SAP applications makes much sense.

MDM can also significantly reduce the burden on SAP systems by providing master data services for master data for non-SAP systems. A robust MDM solution can be much more efficient as a high-transaction, reliable, system of reference for master data.

As an example, consider an e-commerce application that needs to authenticate a customer, and then display the customer's basic information. Rather than have the e-commerce application retrieve the customer data from SAP ECC, e-commerce can retrieve the information from MDM instead, therefore reducing the burden on SAP ECC.

Data quality is one of the key value propositions of adding an MDM solution to an SAP application environment. Without MDM, there is no assurance that standards regarding duplicates or organizational data quality will be adhered to when new master data is entered.

---

[1] Example of the tables for business partners in SAP CRM include BUT000, BUT001, BUT020, BUT021, BUT050, BUT051, and others.

[2] Example of the tables for customer in SAP ERP include KNA1, KNVP, KNVV, KNBK, KNEX, and others.

An estimate is that master data becomes *dirty* at the rate of 2% per month if no data quality enforcement is in place. This is particularly important for SAP applications, because completely removing master data records from SAP after they are entered is often difficult, especially if operational data exists that references the master data (such as orders, invoices, and so on).

The key is to ensure that all master data has been validated and cleansed *before* being entered into the SAP application system. A fully capable MDM solution, such as IBM InfoSphere Master Data Management, has data profiling, cleansing, and onboarding processes built into the solution.

For all of the reasons previously mentioned, an MDM solution *external* to SAP applications, such as SAP CRM or SAP ERP, provides the following benefits to the *enterprise as a whole,* including SAP applications:

- A trusted and complete view of master data across SAP and non-SAP applications
- Easy extensibility
- Central place for business process management (BPM)-based information governance supporting all data stewardship activities, including optimization
- Central place to create and maintain all rules on master data, including but not limited to the following rules:
  - Integrity rules, such as name and address standardization and address verification
  - Matching and survivorship rules
  - Data access rules, such as service and data authorizations
  - Business rules
- Central place for history and audit
- Central place for setting up rules for master data distribution

7.3 Overview of IBM Master Data Management capabilities

IBM MDM has a rich set of functions that cannot all be described in this section. The complete set of functions can be found in 7.2, “Why master data management is important for SAP applications” on page 161. At a high level, IBM MDM provides a ready-to-use, proven, and rich multi-domain data model for domains, such as customer, supplier, product, location, and so on, with support for hierarchy management, relationships, cross-domain relationships, and more.

Several key capabilities of IBM MDM products are as follows:

- Rich set of MDM business services built on service-oriented architecture (SOA) principles.
- Model-driven, rapidly customizable model and services with MDM workbench.
- Optimized for real-time interaction.
- Industry-leading performance and scalability.
- Ready-to-use support for operational, collaborative, and analytical MDM.
- Support for virtual, hybrid, and physical configurations.
- Industry-leading probabilistic matching engine (PME) with a full set of operators to identify duplicates and provide powerful search functions.
- Designed for big data. For example, PME is ported to Apache Hadoop for seamless integration of social personas, enabling social MDM.

- Actionable master data through an event manager supporting life events, time events, and so on.
- Rich security feature set for service and data authorizations, access tokens, and audits.
- History feature capturing data changes on an attribute level. All MDM services are point-in-time enabled.
- BPM-based MDM application for governance and data stewardship enabling seamless collaboration among stewards, extension to mobile channels, and reports for individual steward and stewardship team performance.
- BPM-based MDM authoring processes and MDM Application Toolkit for BPM for quick customization. MDM Application Toolkit for BPM provides business process management-based components that you can use to build MDM business applications. These applications are configured through IBM Business Process Manager.
- Ready-to-use integration with IBM Operational Decision Manager for business rules support.
- Global reach through internalization.
- Accessibility.
- A full-featured reference data management application, IBM InfoSphere MDM Reference Data Management Hub.
- An unstructured text analytics component to enrich customer and product information by analyzing blogs, posts, and more, with customer sentiment, and so on.
- Batch interface for bulk load.
- Integrated across IBM InfoSphere family, with many adapters to third-party solutions.
- Support for multiple deployment options, such as bring your own hardware (BYOH), private cloud, public cloud, and appliances (for example, the IBM PureApplication® System Patterns for InfoSphere MDM).

This rich MDM functionality can be seamlessly integrated with SAP applications, as described in 7.3, “Overview of IBM Master Data Management capabilities” on page 163 and 7.4, “Architecture goals” on page 166.

InfoSphere MDM (see 7.1, “Master data management introduction” on page 160 for more details) is a market-leading MDM platform that both supports all major areas for the adoption of an MDM solution, and can also, as shown in Table 7-1, seamlessly enable many detailed business use cases across many industries.

Table 7-1 IBM InfoSphere MDM is the single solution addressing all use cases

MDM business use case	InfoSphere MDM Enterprise Edition	InfoSphere MDM Advanced Edition	InfoSphere MDM Standard Edition	InfoSphere MDM Collaborative Edition
Advanced catalog management	Y			Y
Asset management	Y			Y
Centrally manage customer preferences	Y	Y		
Customer Information File (CIF) augmentation and replacement	Y	Y		

MDM business use case	InfoSphere MDM Enterprise Edition	InfoSphere MDM Advanced Edition	InfoSphere MDM Standard Edition	InfoSphere MDM Collaborative Edition
Complex product definition	Y			Y
Customer loyalty	Y	Y	Y	
Enterprise master catalog	Y	Y		Y
Federal, State, and Local Citizen Hub	Y	Y	Y	
Global data synchronization	Y			Y
Healthcare payer: Claims, Eligibility, and Member 360°	Y	Y	Y	
Hierarchy management	Y	Y	Y	Y
Improve call center customer service	Y	Y	Y	
Improve campaign marketing effectiveness	Y	Y	Y	Y
Infrastructure rationalization and modernization	Y	Y		
Insurance underwriting	Y	Y	Y	
Law enforcement information exchange	Y	Y	Y	
Mergers and acquisitions	Y	Y	Y	Y
Multichannel commerce	Y	Y	Y	Y
New product introduction	Y			Y
Operational efficiency	Y	Y	Y	Y
Pharmacy exchange	Y	Y	Y	
Parts management	Y			Y
Product factory	Y	Y		Y
Product information management (PIM)	Y			Y
Product bundling	Y			Y
Reference data management	Y	Y		
Risk and compliance	Y	Y	Y	
SOA alignment	Y	Y		Y
Supplier collaboration	Y	Y		Y
Supplier onboarding	Y	Y		Y

7.4 Architecture goals

SAP applications in the SAP *inner ring* and the applications in the non-SAP *outer ring* require access to master data. From an architectural perspective, the following goals need to be achieved by MDM:

- Establish the single version of the “truth” for master data for an enterprise.
- Decouple master data from the consuming applications into a central enterprise-level system.
- Decouple master data from SAP inner ring, because it is required both in the SAP inner ring and in the non-SAP outer ring. Master data is an enterprise resource beyond the SAP environment.
- Provide appropriate interfaces for seamless consumption of master data in a SOA, that is, make MDM a first class SOA citizen.
- Provide appropriate data quality functions to manage master data as a trusted information asset.
- Provide appropriate stewardship functions to govern master data as a trusted information asset.
- Provide appropriate high availability (HA) and disaster recovery (DR) functions. If the MDM capability is unavailable, key business processes might have limited availability, or no availability at all.
- Design MDM with evolution in mind. As relevant business processes change over time, the requirements for master data also change over time, and therefore adaptability is crucial.
- MDM must be flexible to support multiple architectural patterns concurrently, and from an evolutionary phased implementation perspective.

7.5 Architecture overview

InfoSphere Master Data Management delivers enterprise-scale MDM functionality that can serve both the SAP inner ring and the non-SAP outer ring, as shown in Figure 7-2 on page 167. The MDM system manages master data entities, such as customer, supplier, product, and employee, providing master data with the highest degree of quality to all consumers. InfoSphere MDM provides market-leading capabilities to organizations:

- Business services are delivered through intelligent, prepackaged services that can be used to seamlessly integrate MDM into existing business processes and technical architectures.
- Pre-built and extensible data models for master data domains, such as customer, supplier, vendor, employee, location, product, and contract, are optimized for master data management. You can import existing data models or build data models from scratch.
- Ready-to-use MDM application for IBM Business Process Manager for authoring and stewardship of master data.
- Ready-to-use BPM integration with IBM Business Process Manager provides the capabilities necessary to implement policies and coordinate multistep and multirole workflows for data stewardship and data governance.
- Collaborative tasks enable organizations to set up workflows that reflect existing and new business processes, thereby having a system that closely aligns with their business practices.

- InfoSphere MDM Application Toolkit delivers business value rapidly, with governance applications through pre-built blueprints, and widgets for embedding within existing applications.
- Governance and stewardship UIs enable you to inspect and resolve data quality issues in real time, including relationships and hierarchies, and to edit the “golden record.”
- Common PME employs advanced statistical techniques to automatically resolve and manage data quality issues.

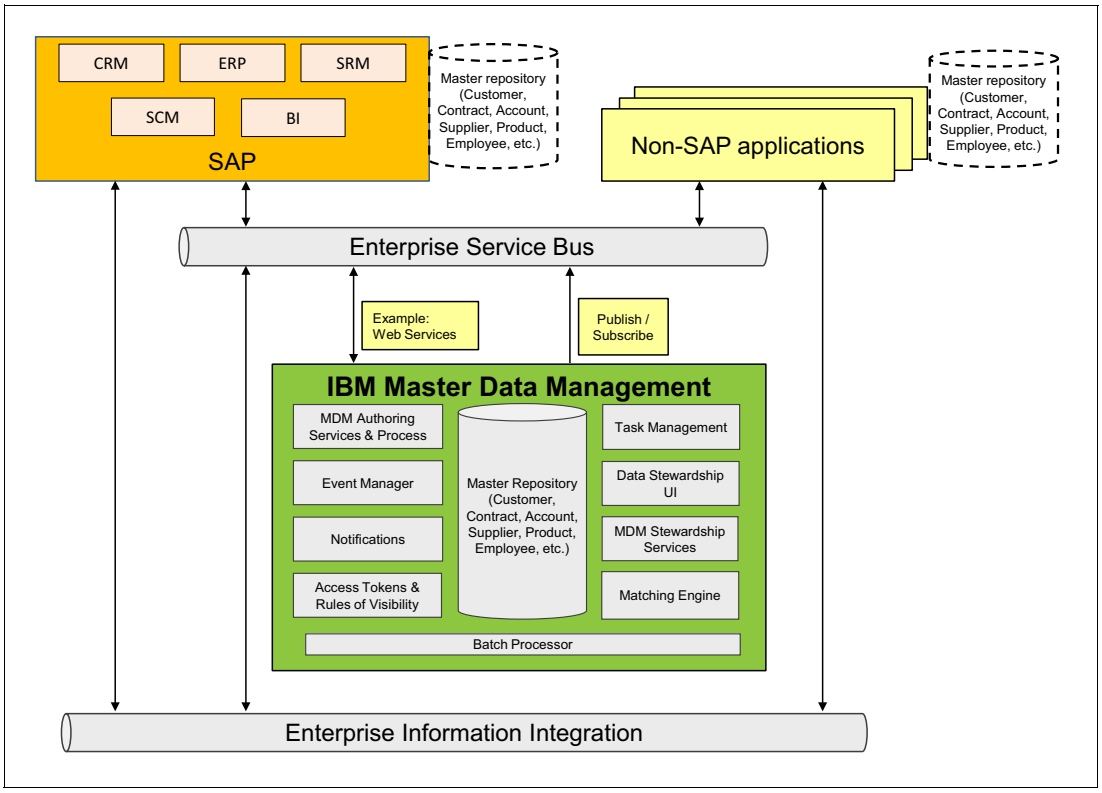

Figure 7-2 IBM Master Data Management for SAP

Figure 7-3 on page 168 shows the components used by the MDM system for efficient integration with all other systems supporting batch and real-time interfaces:

- An enterprise service bus (ESB) component serving both the SAP inner ring and the non-SAP outer ring
- An enterprise information integration component serving both the SAP inner ring and the non-SAP outer ring

In typical implementations, SAP applications hold only a copy of the master data entities (therefore the dotted lines around the master data entities in Figure 7-2), which are centrally managed by the MDM system. The same concept applies to the non-SAP applications in the SAP outer ring.

From both, the ESB and the enterprise information integration components, data exchange must use SAP interfaces such as Business Application Programming Interface (BAPI), Intermediate Document (IDoc), Advanced Business Application Programming (ABAP), or web services. These interfaces are introduced in Chapter 3, “Enterprise integration services for SAP” on page 39.

SAP adapters for IBM Integration Bus (which provides the ESB functionality) and IBM InfoSphere Information Server Pack for SAP Applications, provide certified connectivity for these SAP interfaces. For more information about IBM Integration Bus and IBM InfoSphere Information Server, see Chapter 3, “Enterprise integration services for SAP” on page 39.

7.6 IBM InfoSphere MDM for SAP applications

InfoSphere Master Data Management is an excellent platform for supporting SAP applications. This section takes a closer look at InfoSphere MDM.

As mentioned before, various SAP applications use different data models for the same master data entity. InfoSphere MDM provides a data model that supports the various data models in SAP applications for the same master data entity, such as *customer*, as shown in Figure 7-3 and in Figure 7-4 on page 169. Therefore, InfoSphere MDM provides an excellent foundation for an information technology (IT) environment with SAP and non-SAP applications.

For example, similar to SAP CRM, InfoSphere MDM has separate entities for persons and organizations. InfoSphere MDM supports ready-to-use, multiple-address usage types. This capability provides seamless support to manage the sold-to, ship-to, bill-to, payee, installed-at, and so on, business functions for addresses known from SAP ERP.

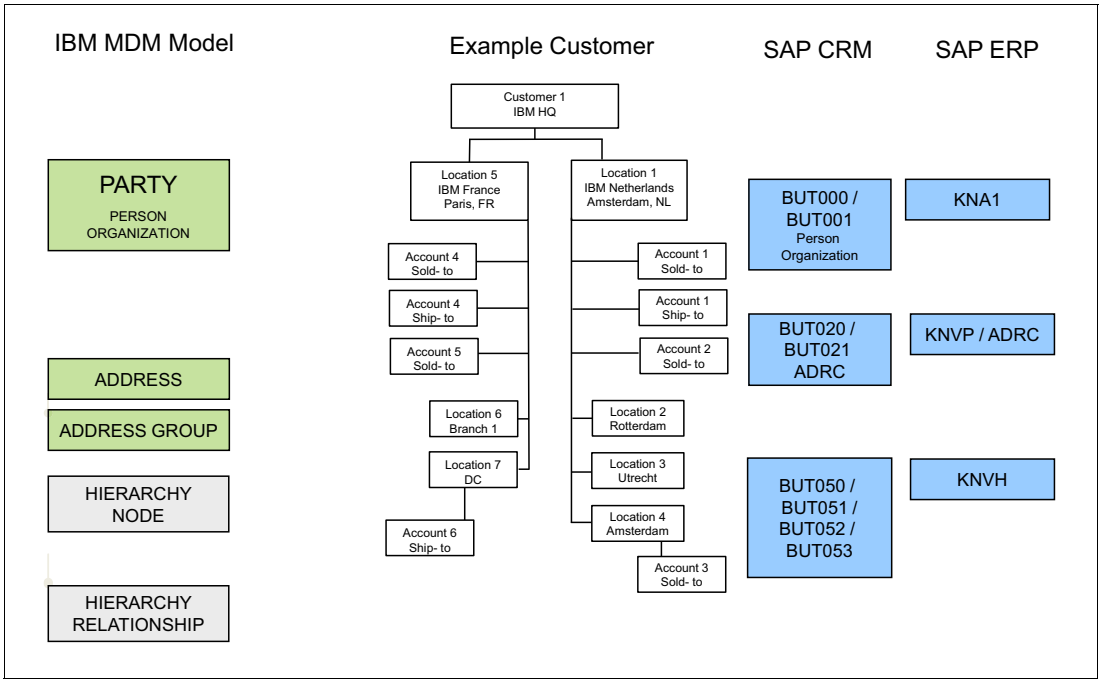

Figure 7-3 Entity-level correlation of InfoSphere MDM data model and SAP application data models

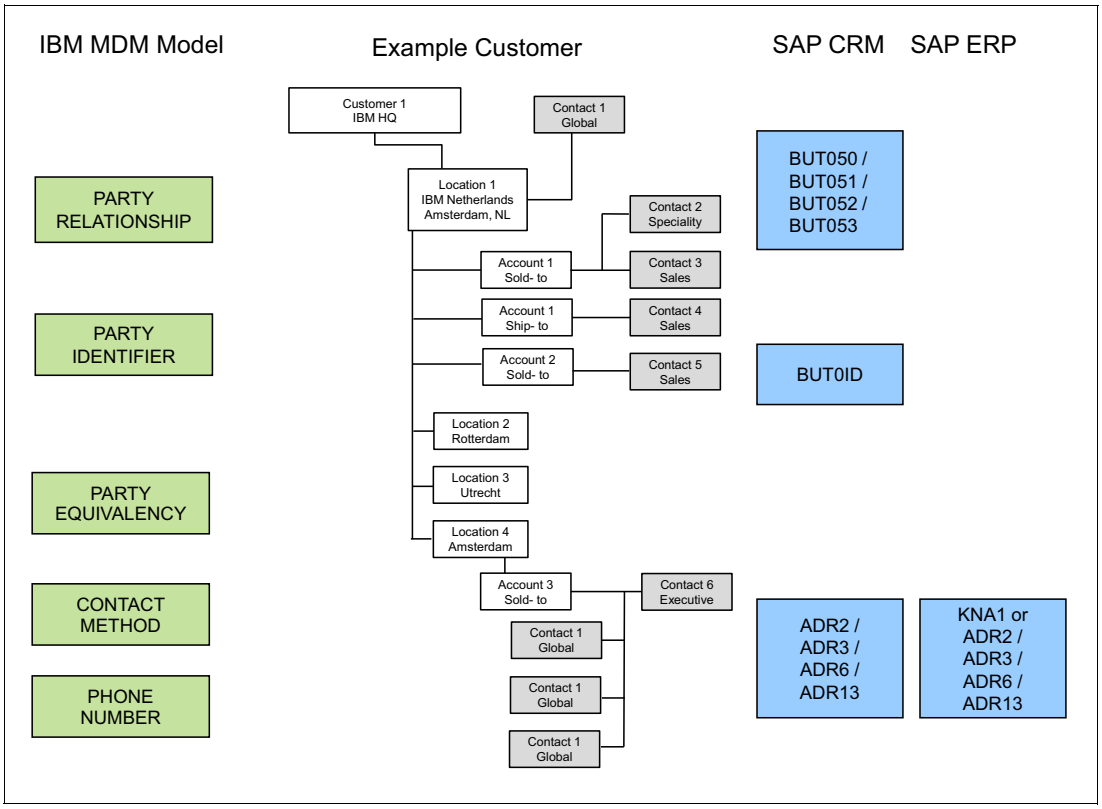

Figure 7-4 High-level entity example between IBM MDM and SAP applications for customer

IBM MDM has many ready-to-use MDM business services that can be consumed through several different protocols, such as web services, Java Message Service (JMS), and so on, making the integration with SAP applications easy. In addition, IBM Integration Bus and InfoSphere Information Server have SAP connectivity on interfaces such as IDoc, BAPI, and so on, for batch, near real-time, and real-time integrations with SAP applications, as described in Chapter 3, “Enterprise integration services for SAP” on page 39.

IBM MDM tools includes IBM Business Process Manager. The MDM application for master data authoring and master data governance supporting stewardship processes is built on IBM Business Process Manager.

IBM Business Process Manager provides ready-to-use integration with the SAP NetWeaver platform. Therefore, if for a specific MDM process, a process integration with SAP solutions is required, this integration can be accomplished seamlessly.

This capability might be needed in the following examples:

- Propagation of master data deletes[3]

  If a master data entity should be deactivated in InfoSphere MDM, there might still be consuming applications, such as SAP applications, using the record. Therefore, the deactivation process must inform the application consuming the record that the record is intended to be deleted. The consuming application must respond back indicating whether deleting the record is okay, and when. In SAP, a master data record cannot be deactivated if pending transactional objects, such as orders, are not finished processing.

  In this case, the following steps must be taken:

  – Set the `block` flag in SAP, preventing new transactional objects to be created for that master data record but enabling the open transactional objects to complete.

  – Set the `deactivation` flag in SAP after the last transactional object has completed.

  – Respond back that deactivation is now OK.

  For this reason, a master data deletion is usually a process that must be orchestrated across multiple systems, including SAP applications. IBM Business Process Manager, part of the IBM MDM portfolio, is an ideal platform for the orchestration.

- Propagation of merge and split operations

  The merge of multiple duplicate master data records, or the split of one master data record into multiple records, has implications in the consuming applications for dependent transactional objects, such as orders. For this reason, these operations require process integration with the consuming applications. IBM Business Process Manager provides an ideal platform.

7.7 Architecture patterns

From a solution architecture perspective, describing all of the details in the architecture overview provided in this chapter is not possible. The IBM MDM Reference Architecture is described in the book *Enterprise Master Data Management: An SOA Approach to Managing Core Information* (see 7.8, “References” on page 187).

This section shows only a few of the key architecture patterns at a high level, to demonstrate some of the options and leading practices. This section is structured as follows:

- 7.7.1, “Master Data Integration” on page 172 introduces how an MDM system is initially built using master data integration.
- 7.7.2, “Master data distribution” on page 175 provides a high-level overview of how master data can be synchronized across the enterprise.
- 7.7.3, “MDM hub patterns and MDM implementation styles” on page 179 provides additional architectural details about various MDM system configuration and implementation patterns.

[3] “Delete” in InfoSphere MDM is a logical deactivation by setting an end date. It is not a physical delete operation.

Before examining the patterns in detail, be sure to understand the basic concepts of system of record, system of reference, core, common, and local, as shown in Figure 7-5.

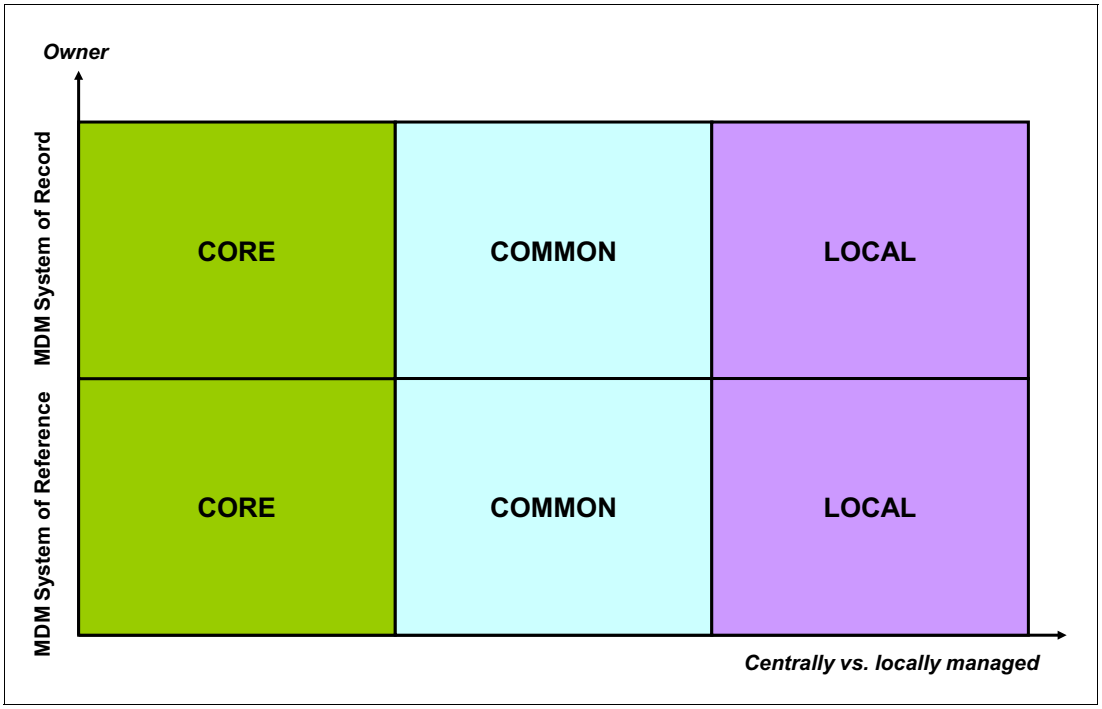

Figure 7-5 MDM concepts: System of record, system of reference, core, common, and local

- Owner. The first dimension is the ownership dimension for an attribute (whether the attribute is created and maintained through the MDM system or outside of MDM).
- System of record. The MDM system is the system of record for an attribute *if it is created and maintained through the MDM system*. Examples could include names, contact details, and so on.
- System of reference. The MDM system is the system of reference for an attribute *if the attribute is created and maintained outside the MDM system*. This could be the case in some implementation styles or for attributes coming from third-party data sources, such as Dun & Bradstreet for the DUNS number, Global Product Code (GPC), and so on.

In addition, the set of master data attributes can be divided into the following categories:

- Core. Attributes in this category are used to uniquely identify a master data entity. These attributes are frequently used in the matching algorithm, and examples include social security number, date of birth, and so on.
- Common. Attributes in this category are used by at least two consuming applications.
- Local. Attributes in this category are relevant only for a single application.

With this definition, only core and common attributes should be managed by the MDM system, because only these attributes have relevance for the enterprise, where application local attributes lack this relevance. Duplicating application local attributes in MDM does not make sense, because no other consumer exists for them anyway. Following this practice provides several benefits:

- Greater flexibility in managing master data in the MDM system, because it does not become a central monolithic system burdened with application local concerns.
- Centralizing only where there is a business case. For example, the customer master data entity in SAP ERP has over 500 attributes. However, in many implementations, only 150 - 200 are relevant beyond SAP ERP, making it a candidate for being managed in the MDM system.
- Reduced cost for central management, because the investment is only on master data attributes that provide value on an enterprise scale.
- Reduced cost using pre-built integration with a flexible system of record and system of reference approach.
- Enabling companies to start small and grow as their needs grow. This approach gives companies the time to mature their *information governance* capabilities and organization in parallel.
- Phased centralization, which is sometimes politically easier because it shows results on every phase, gaining growing trust for the overall MDM strategy.

7.7.1 Master Data Integration

The master data integration (MDI) component of the MDM solution architecture addresses the need to extract and harmonize master data from heterogeneous source systems before loading the data into the MDM system and distributing it using batch processes to other target systems, as shown in Figure 7-6.

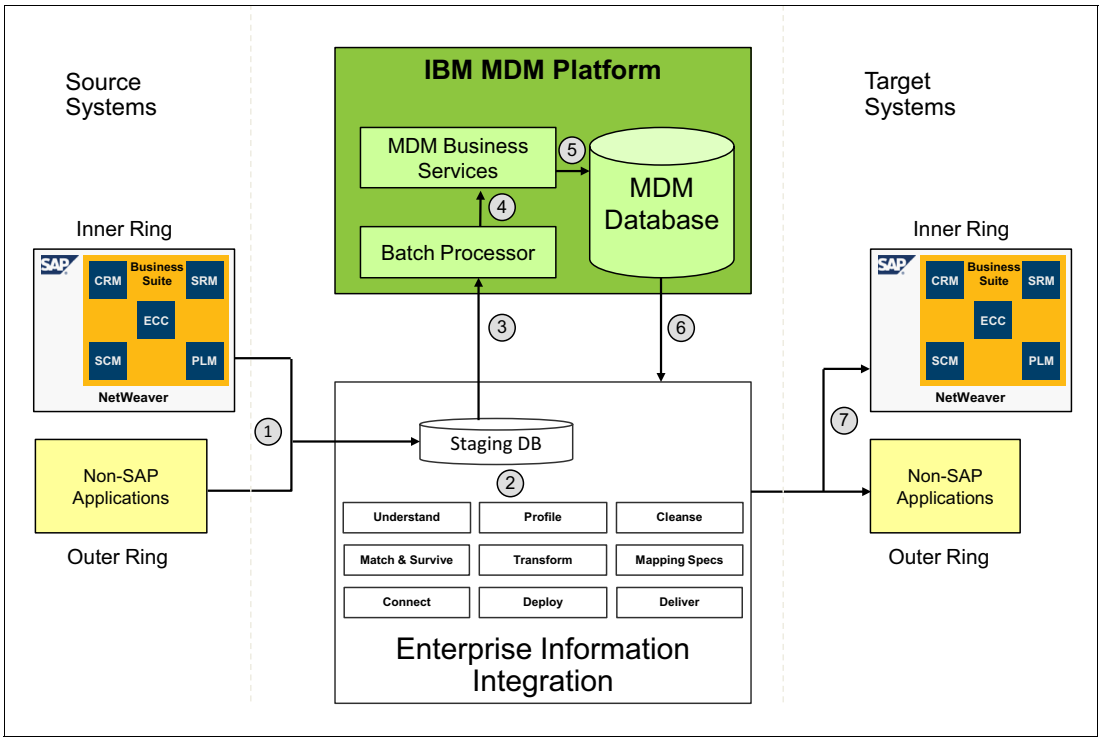

Figure 7-6 Master Data Integration

In some projects, organizations implement new SAP systems in environments that already have existing SAP systems that were consolidated into the new SAP systems. In some cases, several to dozens (and even more than 100 in very large installations) SAP R/3 systems are consolidated into a new SAP ERP system.

In other cases, several SAP ERP systems, originally deployed by geographical area or business unit, are consolidated for process consistency and efficiency into a single SAP ERP system. Therefore, source systems can be SAP and non-SAP applications. From an architecture perspective, all SAP systems are grouped into the SAP inner ring. The non-SAP applications are all other systems that have master data that needs migration to the new SAP systems.

Note that non-SAP applications include third-party data sources, such as Dun & Bradstreet and others, that might be used during the master data harmonization process for enrichment purposes. Target systems include all SAP systems in the SAP inner ring, but can also include non-SAP applications.

The IBM MDM platform has the MDM Authoring Services and Process layer as the entry point, and the MDM database as persistency, as shown in Figure 7-2 on page 167. The Batch Processor is a wrapper around the MDM Authoring Services & Process, orchestrating services and processes in a highly parallel manner for batch loads. A typical situation is the initial load in the first implementation phase, and bulk loads in subsequent implementation phases. The Batch Processor uses as input XML files.

The enterprise information integration component used to perform the data cleansing and harmonization is based on the InfoSphere Information Server product and its components, which work together to achieve business objectives within the information integration domain:

- Understand: IBM InfoSphere Data Architect, IBM InfoSphere Business Glossary, IBM InfoSphere Metadata Workbench, IBM InfoSphere Blueprint Director.
- Profile: IBM InfoSphere Information Analyzer, IBM InfoSphere Discovery.
- Cleanse: IBM InfoSphere QualityStage, IBM InfoSphere DataStage.
- Match and survive: IBM InfoSphere QualityStage.
- Transform: IBM InfoSphere DataStage.
- Mapping specifications: IBM InfoSphere FastTrack.
- Connect: IBM InfoSphere Information Server Pack for SAP Applications (SAP connectors, referred to as *SAP Packs* in this chapter), and so on.
- Deploy: IBM InfoSphere Information Services Director.
- Deliver: IBM InfoSphere Change Data Capture, IBM InfoSphere Federation Server.

IBM InfoSphere Information Server is also the software platform for the conversion architecture. For more information about enterprise integration service, see 3.6, “Initial data load” on page 66.

If the MDM system is deployed parallel to the SAP system, it is common that the InfoSphere Information Server infrastructure is used to support both the conversion architecture and the MDI tasks.

Figure 7-6 on page 172 shows the following key steps in the MDI architecture:

1. Master data from a heterogeneous set of source systems is extracted and stored in a staging database (Staging DB in Figure 7-6 on page 172).

   For SAP source systems, typical extract interfaces include IDoc, BAPI, and ABAP. For SAP sources with the InfoSphere Information Server Pack for SAP Applications V7.0 (SAP Packs v7.0) or newer, two tools based on IBM InfoSphere Data Architect improve the efficiency of this step:

   - IBM InfoSphere Rapid Modeler for SAP Applications (Rapid Modeler)
   - IBM InfoSphere Rapid Generator for SAP Applications (Rapid Generator)

   Rapid Modeler discovers the SAP data models in SAP systems and extracts the SAP data model representing the business objects. Rapid Generator generates the complete extraction logic to read the data from SAP systems and write it into the staging database. The extraction logic is composed of ready-to-run jobs for InfoSphere DataStage.

2. After the master data is extracted into the Staging DB, the following tasks are performed:

   – InfoSphere Information Analyzer is used to *profile* the master data in the Staging DB to understand the data cleansing and harmonization needs.

   – InfoSphere FastTrack is used to *define* the mapping specifications to logically map the various different source models to the MDM target model.

     In addition, physical mapping specifications are defined, which in a first step map the various different source models to a common alignment model where the data cleansing is done. Another specification maps the common alignment model to the MDM target model.

   – InfoSphere DataStage is used to *transform* the master data records from the various different source systems into a common alignment model in which the data cleansing is done.

   – InfoSphere QualityStage and InfoSphere DataStage are used to implement data *cleansing*, such as address standardization (InfoSphere QualityStage) or reference data harmonization (InfoSphere DataStage).

   – Optionally, InfoSphere QualityStage can be used to perform *matching and survivorship* to remove duplicate master data records. This task is optional, because ideally the matching and survivorship logic used for initial load is the same as the one used by the MDM business services when the MDM system is live.

     Therefore, it is usually considered a good practice to use the built-in PME of IBM MDM to detect duplicates and apply appropriate survivorship to them. Step 5 on page 175 provides some more details about this topic.

   – After the data cleansing and harmonization of the master data is complete, transform capabilities of IBM InfoSphere DataStage are used to write the Extensible Markup Language (XML) files for the Batch Processor to process.

3. The Batch Processor reads the XML files containing the clean master data records and starts the MDM business services, once per master data record in the XML files, in a highly parallel manner. The degree of parallelism is configurable to best use the available hardware resources while not overloading the system.

4. The MDM business services started by the Batch Processor process the master data records, and apply all business and data integrity logic to the master data records.

5. If the MDM business services successfully complete their task, they persist the master data record in the MDM database. The MDM business services contain a flag to perform duplicate detection. For large to very large volumes of master data records, it is advisable to disable the duplicate detection during the initial load process, and then schedule an *evergreening* task to perform duplicate detection immediately after load completion but before the MDM system is used.

   The reason for this approach is that duplicate detection is a resource-intensive operation within a single MDM Business Service operation, with measurable effect on throughput of records that can be loaded in a certain time window. A batch duplicate process performed immediately after the load provides for optimal load performance and optimal duplicate detection performance on large volumes of new master data records.

   However, performing duplicate detection and survivorship with the MDM system during step 2 instead has the business benefit that the complete record history is auditable and stored in the MDM system.

6. After the master data is loaded to the MDM system and, if necessary, the evergreening process for duplicate resolution is completed, the enterprise information integration component can be used to extract and transform the master data for the consuming target systems.

7. After the master data is transformed to the format expected by the target system, a bulk load can be performed. For SAP targets, the SAP Packs can be reused to generate the load jobs (similar to the extract jobs).

An implementation leading practice of this architecture is to configure the SAP system for *external key assigned*, and the MDM system creates the primary keys for the SAP system in the correct format. The InfoSphere MDM product includes ready-to-use the necessary data model and services to manage any number of cross-referencing keys, which can be used to maintain these primary keys.

7.7.2 Master data distribution

Harmonizing master data and loading it into an MDM system is only the first step. In a second step, the MDM system must be embedded into the IT environment, which can be done in a variety of ways. The option described in this section is one of several architecture patterns that has been successfully used in many projects.

Figure 7-7 on page 177 shows a conceptual view of embedding the MDM system into the IT environment using the transactional style architecture pattern. For more details about this architecture pattern, see 7.4, “Architecture goals” on page 166. The transactional style pattern establishes the MDM system as the system of record for master data in the enterprise.

Master data creation or change happens through the MDM business services provided by the MDM system. The MDM business services operate on a request-response model. They are transaction-enabled on the service level, which means that they can be used in 2-phase commit protocols, such as WS-Atomic Transaction. An MDM UI is required where master data can be created and maintained. Depending on the business requirements, an MDM UI can mean several different things.

The following list is only for illustration purposes:

- In banking, the MDM UI, in the context of a specific use case, can be the screen of an ATM UI where customers can specify their language preference and other personal settings. When the settings are captured, every time the customer comes back to any ATM machine and swipes the debit or credit card, the ATM UI displays with the customer preferences

previously selected, avoiding the need to select them each time that the customer uses the ATM.

For practical purposes, however, the ready-to-use MDM UI is never displayed on the ATM. The ATM always has its own UI, but with the need to consume MDM services. Therefore, the MDM solution must provide services that have the following attributes:

- Are easy to consume
- Seamlessly scale and perform
- Include high levels of built-in security
- Have the correct level of functionality regarding data model, data quality, and so on, built in

Therefore, these “hidden” attributes of the MDM software are critical aspects to consider when going through an MDM software selection process, rather than placing priority on the UI aspects. Scalability, performance, and the correct functionality are capabilities that require much effort if you have to build them.

- Each time MDM business services are started to create or maintain a master data record, embedded data quality functions are started. Data quality functions can include matching to detect duplicates. Depending on the match result and the business requirements around survivorship rules, it is possible that, in some cases, detected duplicates are merged automatically based on configured survivorship rules.

  In other cases, the detected duplicates are marked as suspects for review by data stewards, because the involved master data records might show some similarity but not enough to warrant an automatic merge. In such a case, data stewards require an MDM UI where they can review the duplicate records that are marked as suspects, and then decide whether they are indeed duplicates.

  Depending on the decision, they might need to manually merge the duplicates, trigger the automatic merge of the records, or mark the records as not being duplicates at all. In any case, the data stewards require an MDM UI enabling them to open master data records, possibly edit them, and enable them for split and merge operations of master data records.

  All of these activities, enabled through the MDM UI, start MDM business services in a real-time fashion. In some cases, the UI used by the data stewards can be a ready-to-use MDM UI. in other cases, it might be a custom-built UI, depending on requirements.

  Also note that the task management and stewardship processes might use workflow capabilities provided by an enterprise BPM platform. In this case, the MDM UI drives the workflow, and individual workflow steps then start the MDM business services in a real-time fashion.

- A customer call center employee might have an intranet portal application where different widgets show various aspects of the customer:

  - The MDM widget provides the ability to quickly retrieve the customer master data record from the MDM system when the customer calls and provides the customer number. This widget starts the MDM business services to pull the customer master data.
  - Other widgets can pull, for example, the contract from a contract management system, such as IBM FileNet. These widgets can also pull, for example, the last five orders from the order management system, and maybe the last 10 interactions that the customer had across all customer touch points of the enterprise from various other systems.

- An enterprise might have an e-commerce channel where customers can register themselves and place orders for various product offerings. In such a scenario, the e-commerce platform starts MDM business services to create new customer master data records during registration, or to update them if a customer is updating address information, contact details, preferences, and so on.

The examples described previously have the following characteristics in common:

- An MDM UI for an enterprise is usually not a single UI. The *MDM UI* might be a dedicated MDM UI in addition to widgets, portlets, and other UI components embedded in other UIs interacting with the MDM system.
- In any case, through the UI, the MDM business services are started, performing real-time transactions on the master data.
- Optionally, the UI might be the entry point to more complex business processes, implemented as workflows on an enterprise BPM component where individual workflow steps start the MDM business services in a real-time fashion.
- Any consumer of master data is receiving only the new master data records, or updates to existing master data records, *after* the MDM system processed it through the MDM business services. An integration interface makes the new or updated master data information available to the consumer.

Figure 7-7 shows an MDM system using the transactional style architecture.

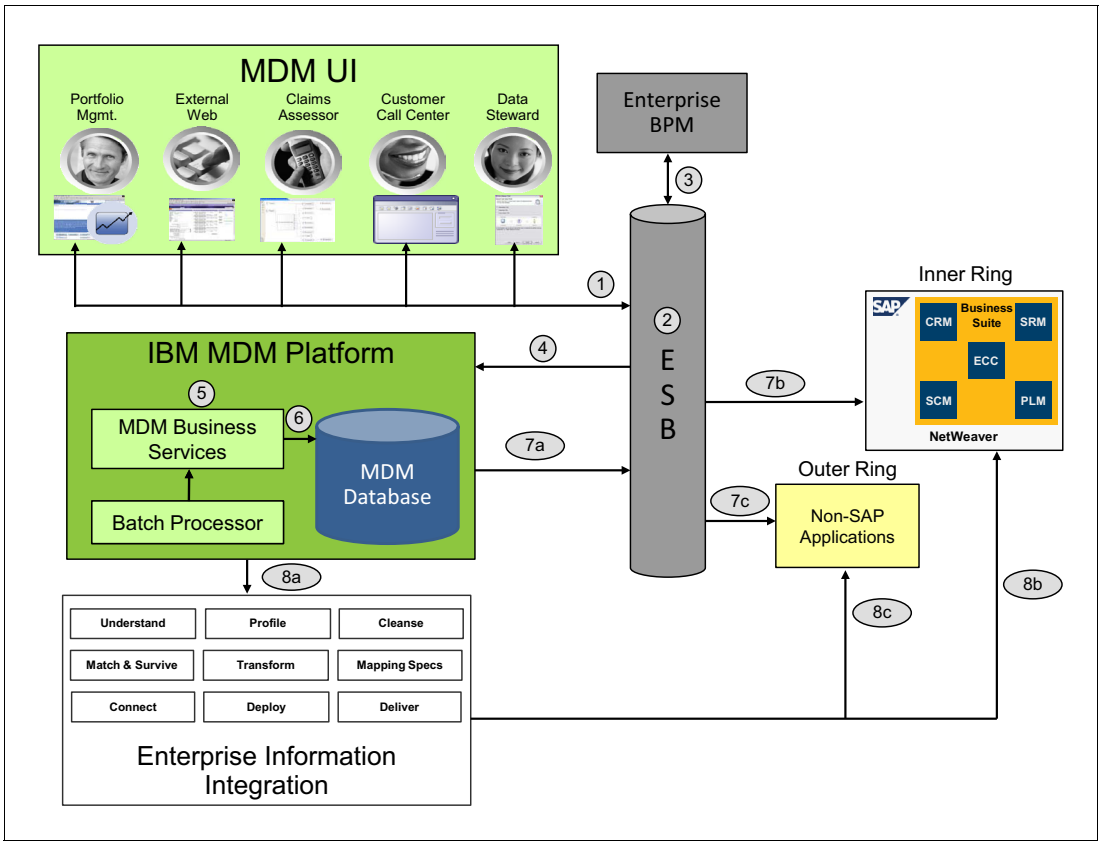

Figure 7-7 MDM system using transactional style architecture

Figure 7-7 shows the following key steps in the architecture:

1. A user uses an MDM UI component to perform search, create, or maintenance functions on master data.
2. Through the ESB, the MDM business services are started for real-time, transactional operations on service level.
3. Optionally, if the MDM UI is driving business processes implemented as a workflow on an enterprise BPM component, workflow steps request the MDM business services.

4. The ESB routes the MDM business service request to the MDM system.
5. When the MDM system receives the MDM business services request, it starts processing.
6. If the MDM system completes the processing successfully, it writes the appropriate master data updates to the MDM database, and returns a success status to the service requester through the ESB. Otherwise it returns an error.
7. Upon successful completion of the MDM business services, a notification to the consumer is triggered:
   a. The MDM system has a notification framework that can be used to publish changes on master data to subscribed consumers, using publish/subscribe patterns on the ESB.
   b. Because many large enterprises implement the ESB using IBM Integration Bus, which is the industry-leading ESB platform for consuming SAP applications, (near) real-time integration can be implemented by using IBM Integration Bus with IBM WebSphere Adapter for SAP Software. WebSphere Adapter for SAP Software provides seamless integration with standard SAP interfaces, such as IDoc, BAPI, and so on.
   c. For master-data-consuming, non-SAP applications, depending on the application type, other connectors of IBM Integration Bus (previously known as IBM WebSphere Message Broker) can be used for (near) real-time integration.
8. Some consumers of master data, such as analytical applications, might not require real-time or near real-time updates on master data. For these consumers, periodic batch updates might be sufficient.

   These batch updates can be implemented using the enterprise information integration component:
   a. A periodic batch is extracted from the MDM system based on the scheduled frequency (hourly, daily, and so on).
   b. For SAP systems with only periodic update needs, the enterprise information integration component transforms the periodic batch to the SAP data model, and uses the SAP Packs to load it to the consuming SAP application.
   c. For non-SAP applications with only periodic update needs, the enterprise information integration component transforms the periodic batch to the appropriate target model, and uses other provided connectors to load it into the target system.

From an MDM perspective, the ESB is used to achieve at a minimum the following objectives:

- Loose coupling of the MDM system and consuming applications (SAP applications and Non-SAP Applications).

  The primary purpose for this objective is to avoid point-to-point integration between the MDM system and consuming applications. Because the MDM system evolves over time (for example, a change in business process can drive changes in the data model and services interface of MDM), a single change can break all point-to-point integrations between MDM and the consuming applications.

  Therefore, use the publish/subscribe capabilities offered by the ESB, so that MDM can notify subscribed consumers of changes in master data as needed. The publish/subscribe approach for loose coupling is further enhanced if you use the ability to create a canonical data model with appropriate governance in the ESB.

  For example, three versions exist concurrently. If a new one is added, the oldest is withdrawn from service. This approach enables consumers to stay on a certain version for a longer period of time. This avoids interface rework each time a new consumer, or new requirements for an existing consumer, create a change to the canonical data model.

  Also, if an update to the interface becomes necessary because the current version used by a consuming application is the oldest of the concurrently supported versions, an

upgrade to the *newest* canonical data model is possible. Skipping version changes in between, therefore reducing the number of interface changes, reduces operational integration cost.

- For seamless connectivity to consuming applications, the application connectors offered by an ESB are used. InfoSphere MDM can, for example, seamlessly synchronize master information to subscribed SAP applications using an IBM Integration Bus-based ESB system and the WebSphere Adapter for SAP Software, on interfaces such as BAPI and IDoc.

7.7.3 MDM hub patterns and MDM implementation styles

For implementing an MDM system, no “one size fits all” model exists. For example, in healthcare, the requirement can be to implement an MDM solution for patient master data. The constituents in healthcare are the patients, health service providers, such as doctors and hospitals, and health insurers.

A healthcare insurer gets patient information from the patient when the patient signs the health insurance contract. The healthcare insurer also receives patient records as part of the bills from doctors and hospitals, where the information for the same patient might have been captured differently.

However, the healthcare insurer has no control over how the patient information has been captured at a doctor’s office or in a hospital. Consequently, for the same patient, multiple patient records exist, some of them sourced outside of the healthcare insurer’s IT environment.

An MDM system managing patient information in such a scenario must be able to link the patient records from the various sources correctly. This is a typical scenario for virtual hub configurations for MDM.

The banking industry provides a different example, where customer master data is centrally created and maintained through the MDM services of a central MDM system. Master data authoring in such a scenario always operates on the same master data record using MDM services. The master data record is accessed from online and mobile banking channels, core banking applications, ATMs, and so on. The physical hub configuration for the MDM system is usually the best fit for this type of solution scenario.

The examples described previously show that various configuration requirements exist for MDM systems to address specific needs. The following sections describe some of the options.

Virtual hub, hybrid hub, and physical hub patterns

As illustrated by the previous industry examples, different MDM system configurations are required:

- Virtual hub
- Physical hub
- Hybrid hub

Each configuration type has the following high-level key characteristics:

- Virtual hub

  Master data authoring is done through LOB applications, portals, and so on. As part of this process, a thin slice of the master data attributes for a record is sent to the virtual hub and persisted there *as-is*. Because it is done from multiple systems, it is likely that, for the

same entity, multiple records exist in the MDM system coming from different sources. The source records are known as *silver* records.

A virtual hub computes the *golden* record in real-time from the various source silver records whenever an access to the golden record is needed. However, the golden record is not physically persisted.

- Physical hub

  All relevant core and common master data attributes are stored in the MDM system. Core attributes are those required for unique identification. Common attributes are those with at least two consumers.

  Master data authoring is frequently done in the MDM system.

  All changes are done against golden records only.

- Hybrid hub

  Conceptually, the MDM system has a virtual (silver record) and a physical (golden record) storage area.

  In this configuration, a new record is created through the virtual side of the hybrid hub and, as soon as the golden record is computed, it is asynchronously moved to the physical side of the hybrid hub and continuously updated from the virtual side as needed.

  In addition, on the physical side of the hybrid hub, MDM services can be used to extend the thin golden records coming in from the virtual side as necessary.

In practice, an advisable approach is sometimes to start first with a virtual hub configuration and gradually move to a physical hub configuration as business needs, in-house skills, and maturity grow for the MDM system.

MDM implementation styles

In addition to deciding on the configuration option for the MDM system (virtual hub, physical hub, or hybrid hub), many architecture patterns are also available that can be followed to embed the MDM system in the IT landscape.

Examples include consolidation, registry, co-existence, and transactional style patterns that are described in the following sections:

- 7.2, “Why master data management is important for SAP applications” on page 161
- 7.4, “Architecture goals” on page 166

A high-level summary of the MDM implementation styles is described in the following list:

- Consolidation style (see Figure 7-8)

  In this implementation style the MDM system is a *system of reference* because the authoring and maintenance of master data happens outside of the MDM system. The master data is physically materialized in the MDM system, which is a consolidation point for some consuming applications. Consuming applications are typically analytical systems, such as data warehouses or systems for external consumption, such as product master data in manufacturing, which needs to be consumed by many external retailers.

  The benefits and disadvantages of the consolidation style are as follows:

  – Benefits. The consolidation style has no effect on source applications, because business users can create and maintain master data as they were used to doing before. It therefore enables a low-risk start on the MDM journey, and *full MDM benefits for downstream consumers* by consuming high-quality, consistent master data from the MDM system. For the consumers, they can now get all master data from a single place, rather than having to integrate with multiple sourcing systems.

  – Disadvantages. The *source applications do not benefit from the MDM system*, which means that the master data quality in sources remains low, with possibly negative effect on operational processes.

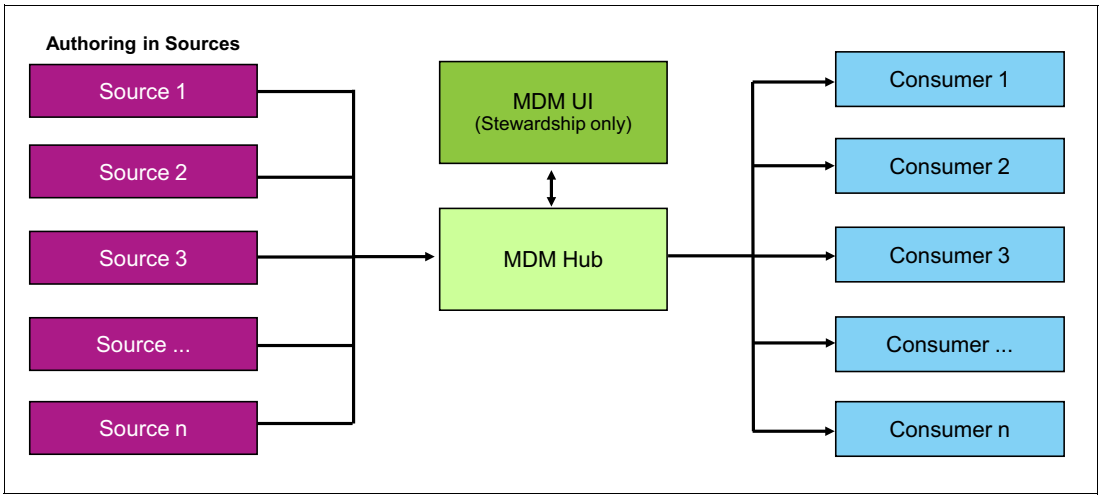

Figure 7-8 Consolidation style

- Registry style (see Figure 7-9)

  The MDM hub for this implementation style is a hybrid of system of reference and system of record.

  The MDM hub is a *system of record for a small subset of physically materialized master data attributes* primarily used to enforce uniqueness.

  The MDM hub is a *system of reference for the majority of the master data attributes* by cross-linking to the sourcing applications that are the system of record for them.

  Because of its characteristics, a complete master data record is dynamically created at run time, often read-only, and the authoring remains distributed.

  The primary use case for this implementation style is the centralized assignment of enterprise IDs for master data entities, and central deduplication.

  The benefits and disadvantages of the registry style are as follows:

  – Benefits. This implementation style has a low effect on source applications, because sourcing is still mostly done through them. However, it *establishes a corporate-wide numbering system using enterprise IDs* that is, for example, useful for downstream consumers, such as data warehouses and other analytical systems.

  – Disadvantages. The major disadvantage of this implementation style is that *data quality cannot be centrally managed*, because different sourcing locations remain, and therefore inconsistencies across sources and consumers remain. Also, complete golden records are sometimes available through only a federated query approach.

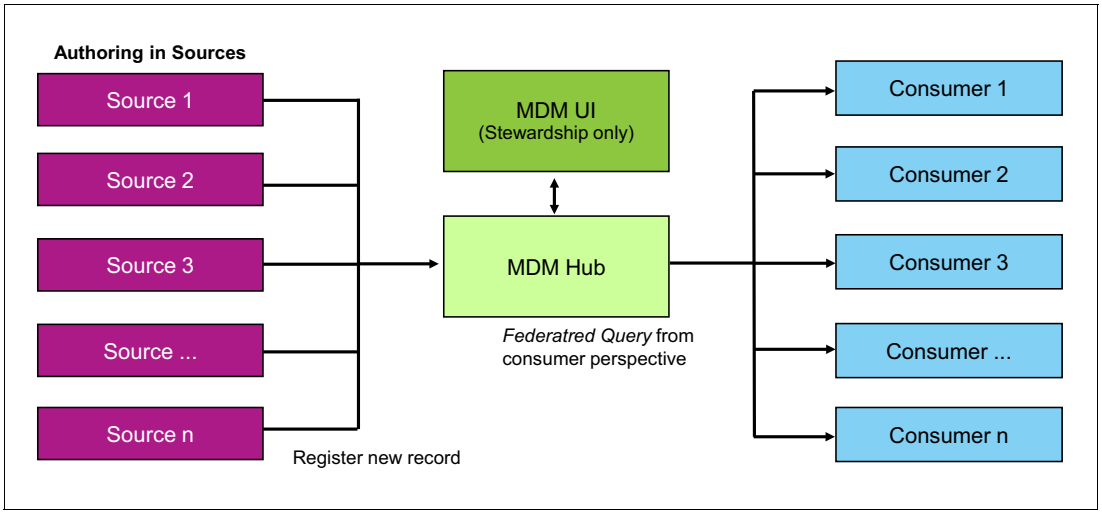

Figure 7-9 Registry style

- Co-existence style (see Figure 7-10)

  With this implementation style, the MDM Hub is a *hybrid of a system of reference and system of record*, because authoring is done through the MDM Hub and in other applications. The fact that authoring “co-exists” across the MDM system and other sourcing applications is part of what gave this implementation style its name.

  In this implementation style, *master data is fully physically materialized in MDM*, which distinguishes it from the Registry Style, for example, which has only a thin slice in the MDM system. The primary use of this implementation style is for data harmonization across systems and central reference for downstream consumers.

  The benefits and disadvantages of the co-existence style are as follows:

  – Benefits. *Full MDM benefits are available for downstream consumers* by consuming high-quality, consistent master data from the MDM hub. An additional benefit of this style is that the source applications get a write-back from the MDM system, improving their data quality.

  – Disadvantages. *It is difficult to implement bidirectional data synchronization* to keep in sync the applications that retain create and maintain functionality with the MDM system. What makes bidirectional data synchronization difficult to implement is the fact that a master data record can be changed in the MDM system *and* in the applications outside of the MDM system.

    For such updates, conflict resolution is needed, which is complex and therefore error-prone in practice to implement. If the bidirectional data synchronization is not done in real-time, master data in the MDM hub can be stale, causing master data inconsistencies.

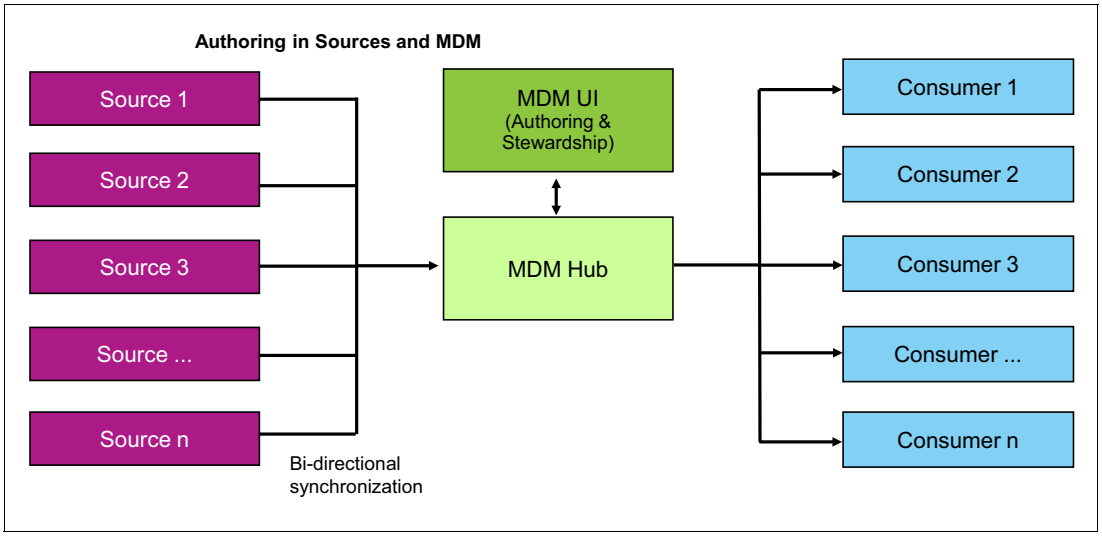

Figure 7-10 Co-existence style

- Transactional style (see Figure 7-11)

  In this implementation style, *the MDM system is the system of record*, and all master data authoring takes place only through services provided by the MDM system. All master data attributes relevant to the enterprise are materialized in the MDM system. The primary use cases for the transactional style are scenarios requiring real-time transactional access to master data, and central reference for downstream users, such as analytical systems.

  The benefits and disadvantages of the transactional style are as follows:

  - Benefits. This implementation style provides *real-time, consistent access to master data*. The single authoring place simplifies information governance and stewardship because of its centralized structure for maintenance of master data. It provides *full MDM benefits to all consumers* and is the one-stop shop for all master data needs from the consumers' perspective.
  - Disadvantages. The one disadvantage of the transactional style is that it might be *more intrusive for some source systems*, because it can require changes of master data processes and UIs to consume the MDM services from the MDM system.

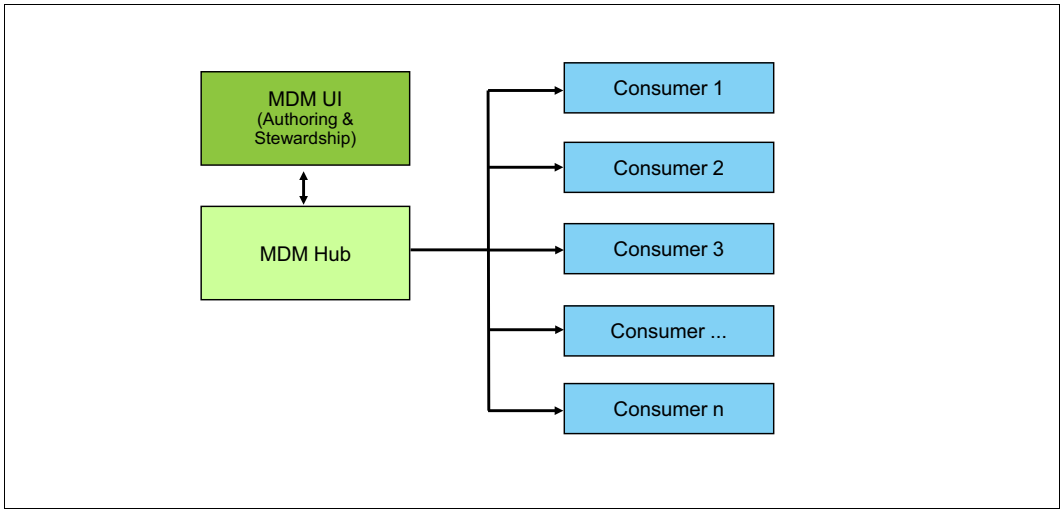

Figure 7-11 Transactional style

In practice, consider the following two aspects:

- MDM environments mature over time, starting with a consolidation style architecture, which is relatively simple from an integration perspective (often a nightly batch). Then the MDM system matures incrementally toward a transactional style architecture. Transactional style is more sophisticated from an integration perspective, because existing master data processes must be standardized, and the MDM services must be consumed.
- Considering business requirements, technical, and cost reasons often leads to MDM systems that are hybrid architectures. For each application, a decision is made on which implementation style to use to integrate the application with the MDM system.

7.7.4 Selecting MDM hub and MDM implementation styles for environments with SAP applications

DM hub configurations and MDM Implementations styles are introduced in “Virtual hub, hybrid hub, and physical hub patterns” on page 179 and “MDM implementation styles” on page 180.

Based on experience of the authors, the following options are the most common in support of SAP applications:

- Physical hub and consolidation style pattern
- Physical hub and transactional style pattern
- Virtual hub and registry style pattern

Although other options have been successfully deployed under certain constraints, the previously listed combinations are the most commonly found in practice for the integration of MDM systems and SAP applications, because of their ease of use.

From an enterprise architecture perspective, the following architecture principles for the MDM solution have been established:

- The MDM system should be installed with the physical hub configuration.
- The MDM implementation style should be transactional style wherever possible, so that authoring and stewardship take place through the consumption of MDM services. Consequently, most of the attributes are *system of record* attributes in MDM.
- Exceptions using the system of record and system of reference classification for attributes are permitted on a per-attribute basis through an architecture governance process, if a sound justification can be provided.
- Applications having a need for master data can complete the following actions:
  - Receive master data updates through a publish/subscribe pattern used on the ESB.
  - Receive periodic batch updates through the enterprise information integration platform.
  - Retrieve or update master data by starting the MDM services.

Figure 7-12 provides a final example to complement the information about MDM architecture options and implementation styles described in this chapter.

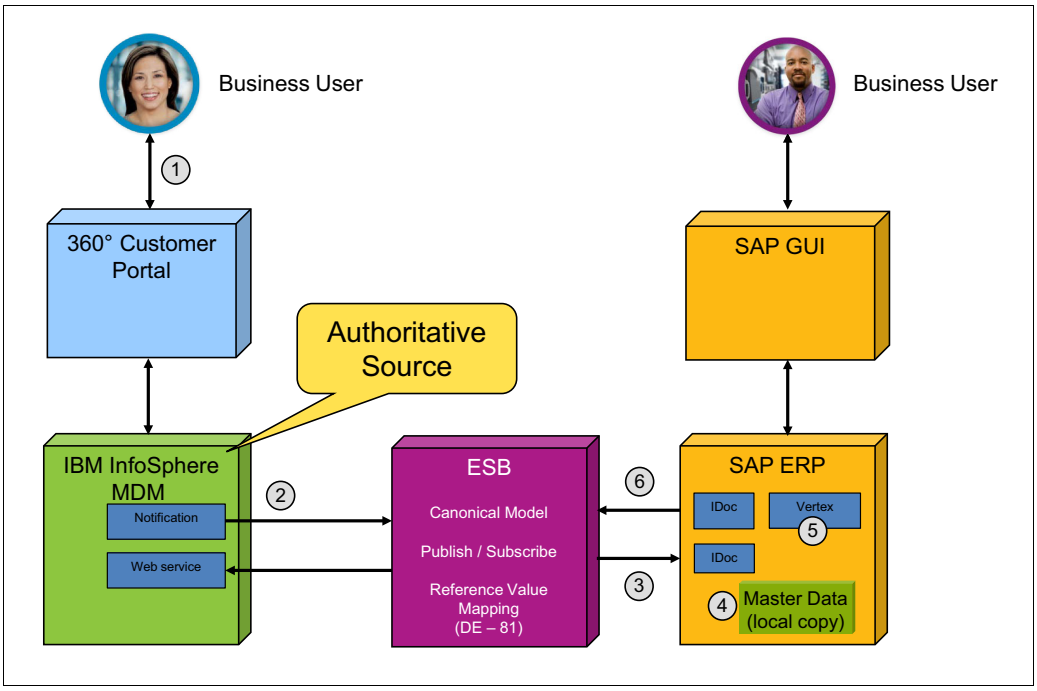

Figure 7-12 Hybrid MDM architecture

The example in Figure 7-12 describes the following flow:

1. A business user creates a new customer record through an intranet customer portal, starting MDM services from IBM Master Data Management.
2. When the customer record is created, the MDM notification feature publishes a message to the ESB in the Canonical Model.
3. An SAP ERP application subscribes to the ESB, which is implemented using the IBM Integration Bus with WebSphere Adapter for SAP Software. Because of this subscription, the new record is sent to the SAP application using an instance of the DEBMAS IDoc structure for customer from the ESB. Before sending the update to the SAP application, the ESB also transcodes the reference values from MDM to the reference values of the SAP application as needed. For example, MDM might store the country for Germany as `81` while SAP ERP has `DE` as the corresponding value.
4. When the SAP ERP application receives the new master data record through an instance of the DEBMAS IDoc, the SAP ERP system persists the master data record locally.
5. VERTEX is an SAP ERP extension that can enrich customer master data with tax-related information. Because the integration between SAP ERP and VERTEX is ready to use, building a custom integration between MDM and VERTEX is not cost-effective.

   After VERTEX enriches the customer record with tax information, it should be distributed through MDM to other consumers also. The objective is to replicate the subset of tax information that is relevant back to the MDM system. For this reason, SAP ERP has prepackaged capabilities to automatically send notifications to other systems, for example, change pointers producing IDOc instances.
6. The ESB “listens” for the SAP notifications and, after relevant tax information updates are received, it routes them back to the MDM system, applying the reverse transcoding operations on reference values by starting the appropriate MDM services.

This example shows that the architecture has the following benefits:

- Where available, it uses ready-to-use integration to reduce costs, for example VERTEX and SAP ERP.
- It provides architectural flexibility. All attributes are system of record attributes except the tax information attributes, which are system of reference in MDM (created and maintained in SAP, but distributed through MDM to all applications that need it as part of the master data synchronization).

More information: For more details, see 7.5, “Architecture overview” on page 166 and 7.6, “IBM InfoSphere MDM for SAP applications” on page 168.

7.8 References

A comprehensive and in-depth description of the MDM Reference Architecture, MDM-related architecture patterns, and deployment leading practices can be found in the following resources:

- IBM InfoSphere Master Data Management

  http://www.ibm.com/software/data/master-data-management/
- IBM InfoSphere MDM version 11.0 documentation in the IBM Knowledge Center

  http://pic.dhe.ibm.com/infocenter/mdm/v11r0/index.jsp
- Dreibelbis, A., Hechler, E., Milman, I., Oberhofer, M., van Run, P., Wolfson, D.: *Enterprise Master Data Management: An SOA Approach to Managing Core Information*, IBM Press

  http://www.ibmpressbooks.com/store/enterprise-master-data-management-an-soa-approach-to-9780132366250
- Hechler, E., Oberhofer, M., van Run, P.: *Implementing a Transaction Hub MDM pattern using IBM InfoSphere Master Data Management Server*

  http://www.ibm.com/developerworks/data/library/techarticle/dm-0803oberhofer/
- Grasselt, M., Nelke, S., Schoen, H.: *Integrating MDM Server with Enterprise Information Systems using SAP as an example, Part 1: Delivering customer records to SAP*

  http://www.ibm.com/developerworks/data/tutorials/dm-1108integratingmdmserver1/
- Grasselt, M., Nelke, S., Schoen, H.: *Integrating MDM Server with Enterprise Information Systems using SAP as an example, Part 2: Enriching customer records with SAP-specific information*

  http://www.ibm.com/developerworks/data/library/techarticle/dm-1307integratingmdmserver2/index.html?ca=dat-

8

Enterprise Content Management for SAP

IBM Enterprise Content Management (ECM) portfolio software enables the largest organizations in the world to make better decisions, faster. By gaining control of structured and unstructured information, organizations can access, collaborate, and influence business decisions in new ways, making content a first-class source of insight.

With industry-specific IBM ECM solutions, organizations can capture, manage, and share content throughout the information lifecycle, helping to ensure compliance, reduce costs, and maximize productivity.

The IBM ECM portfolio includes a wide array of capabilities that integrate with existing systems to help organizations maximize the value of information, including IBM DataCap document imaging and capture, social content management, advanced case management, IBM Information Lifecycle Governance solutions, and IBM Content Analytics with Enterprise Search.

This chapter describes the business goals that can be achieved by extending the SAP archiving infrastructure with IBM Enterprise Content Management portfolio solutions. It explains how the IBM ECM solutions support and extend the existing SAP data and document archive functions. It provides examples for enhancing the implementation of core end-to-end business process solutions by seamlessly integrating SAP and IBM ECM components.

This chapter includes the following topics:

- 8.1, “Enterprise content management business goals” on page 190
- 8.2, “ECM for SAP use cases and solution architecture” on page 196
- 8.3, “Business process enhancements through ECM for SAP solutions” on page 207
- 8.4, “Data governance: Managing growth and compliance” on page 221
- 8.5, “References” on page 230

© Copyright IBM Corp. 2014, 2015. All rights reserved.

8.1 Enterprise content management business goals

This section has the following information:

- Explains the drivers that motivate an organization-wide approach to content management. This approach enables a unified view and management of all business content.
- Describes the importance of a single *source of truth* for all of the information spread across an organization. Challenges are associated with such a requirement, such as managing the volume of information and fulfilling legal requirements for retention and disposal of information.
- Describes the types of information in an organization:
  - Information in databases, typically referred to as *structured information*
  - The vast variety of all kinds of documents that exist throughout a company, both in paper and electronic form, known as *unstructured information*
- Shows how such a holistic approach naturally extends into the information contained in an SAP system, in the form of the raw SAP database content, and in the collection of documents associated with this content as attachments, reports, and other documents in textual or image form.
- Outlines how an integrated and federated solution based on IBM Enterprise Content Management portfolio software supports the goals of different stakeholders in an organization, such as line-of-business decision makers, service-providing information technology (IT) departments, and legal counsels. The solution, based on IBM Enterprise Content Management portfolio software, provides a common approach, a common collection of governing rules, and, as a result, a common view on all relevant business content.

8.1.1 Information lifecycle management: More than just archiving[1]

Any decision-making process in a modern organization has to rely on the availability, timeliness, and accuracy of current business information.

This requirement is associated to several challenges:

- The amount of information that has to be collected, maintained, surveyed, evaluated, and managed grows at a massive rate.
- The vast variety of information resources, information duplication, and uncontrolled versioning makes an efficient discovery process extremely complex.
- An increasing number of rules and regulations influence how information has to be retained or disposed.

Managing the overall growth of data

People who are not directly involved in supporting an IT infrastructure are often surprised to hear that during the last five years, in a typical organization, there has been an average 50 percent year-over-year growth in the volume of data. However, that statistic is not surprising to IT managers, who struggle to find a way to support rapidly growing volumes of data with an often flat or only slightly increasing budget.

[1] Unless stated otherwise, when the term *archiving* is used in the context of this chapter, it refers to the operation that stores content in a repository outside of the SAP system, where it is still active, and immediately accessible. Only in the case of *data archiving*, which is described later in this chapter, does archiving refer to taking inactive content out of the database for performance maintenance reasons.

Also, no surprise to IT managers is that much of the data under management is *debris* that is outdated and duplicated many times across multiple systems, with no real value to the business. Such over-retention results in direct and indirect costs on several levels:

- Overspending on a more complex IT environment:
  - More IT resources required to support larger systems
  - More storage and higher demands to maintain predefined system performance levels
- Higher costs for e-discovery processing. With larger amounts of information, much of it not valuable, processing any requests for information takes longer and results in higher review fees.
- More legal risk inherent with a larger information set.

The proliferation of collaboration products and social media platforms has added both to the diversity and to the volume of data that needs to be growth-managed, and emphasizes the need for an organization-wide solution to manage enterprise content.

System performance

Constantly growing amounts of data not only increase the cost of storage required to keep the data, it also has a detrimental effect on system performance of business-critical applications (apps). High volume accumulation of transactional data in a high-performance business application usually results in a deterioration of application performance, jeopardizing service level agreements (SLAs) for guaranteed response times.

Storage costs use up the IT budget of an organization

If storage of data is not actively maintained, and permanent growth of data without controlled disposal of obsolete data is permitted, an increasing percentage (not just amount) of the IT budget is spent on storage alone. The perception prevails that *storage is getting cheaper*, but the percentage of a typical IT budget spent on storage grows. Identifying and actively managing information debris can result in storage savings of up to 30 - 50%.

The IBM ECM product portfolio offers a family of content collection and archiving products designed to help curtail over-retention by attacking the problem at the source, and minimizing the amount of unnecessary information (the debris) that IT and other stakeholders must manage.

Data governance and discovery

Compliance and corporate information oversight demands require proper governance and defensible disposition of practically all types of content, including archived structured data.

Information discovery is essential

Access to all business-relevant information within the organization needs to be seamless and easy. The search, classification, and categorizing capabilities of IBM ECM portfolio products provide targeted results from the information sets managed by the content repositories.

A common client infrastructure based on IBM Content Navigator as the integration platform acts as a state-of-the-art front end to all discovery services. It is built on open standards using Hypertext Markup Language version 5 (HTML5) and Dojo components. This approach enables delivery of information to the correct client at the correct time.

An enterprise-wide approach is needed

The implementation of an organization-wide content management strategy with federated, deduplicated content requires that all players in the company's decision process network are involved. They must work together on defining a common strategy and a common understanding of the rules that govern the management and organization of all of the company's information throughout its entire lifecycle.

Information "islands" or silos must be abandoned. The uncontrolled evolution of diverging, and often contradictory, rules for governing content must become an integrated lifecycle model agreed upon by all stakeholders.

The information lifecycle and its governance

This section reflects on the lifecycle of information as it occurs in a typical organization, and provides observations on the constraints and challenges that come with it.

Inception of information

For most organizations, the delivery of a product or service depends on the exchange of documents that are part of the record for all transactions. How efficiently organizations manage documents can have a huge effect on the quality of the experience for their customers, patients, students, or constituents.

For all documents that have not been created electronically from the beginning, moving away from the handling of paper as soon as possible in the lifecycle is the first step in this direction. The need for a sophisticated capture solution to digitize the document content is a key requirement.

Smarter document capture uses technologies that convert documents to searchable images, automate data entry, identify documents, check data quality, and format data for adequate use by business systems. By automating labor-intensive, error-prone manual processes, IBM Datacap can accelerate document processing capacity, improve the quality of the processing results with significantly less manual intervention, and reduce cost.

More importantly, IBM Datacap can remove many of the obstacles that degrade service quality, to help organizations create deeper engagements with their customers.

Content creation, revision, and approval workflows

In addition to digitized paper documents, electronically created documents encompass an increasingly large portion of the unstructured content that has to be managed. The creation process, and the revision with proper version control and collaboration between different actors that play a role in the creation of the document, can be modeled in the information lifecycle under the control of IBM ECM software.

In addition to their state-of-the-art storage organization, search, and retrieval capabilities, IBM content repositories provide powerful workflow engines that can be used to model business workflows, such as approval processes, based on an electronic document flow through the organization. Security and audit functionalities ensure that all business decisions follow a well-defined protocol in an auditable way, compliant with the governing rules.

These workflow capabilities can, and in many cases must, be extended and intertwined with the business workflows modeled in the SAP infrastructure of the organization. Experience shows that organizations have business requirements for all of these processes inside and outside of their SAP systems. A critical step is to provide the correct information to the correct client at the correct time. Implementing fully integrated workflows is crucial to satisfy this requirement.

IBM Defensible Disposal

The IBM Defensible Disposal solution applies to the entire organization, which must include content that is also related to SAP systems. As mentioned before, treating SAP structured and unstructured content in a siloed solution counteracts the proposed integrated information lifecycle model, compromises the completeness and accuracy of discovery operations, and raises the cost of disposal under a common retention rule governance.

Also as indicated earlier, not all information that is created and stored in a content repository retains its business value indefinitely. Business transactions have a natural lifecycle. The information associated with the business transactions should be retained according to the retention rules, and subjected to disposal processes to keep data growth under strict control.

It is mandatory, under the legal rules defining compliance and auditability of the business processes, that the disposal process follows a common strategy across all divisions and departments that are involved in the handling of business records. Simply, the decision about when certain content can be disposed of cannot default to the IT department, which might make such a decision indiscriminately, based on the age of documents.

The IT department might be unaware that certain types of content fall under different and often complex rules for their retention and disposal. An example of such complexity is the rules governing human resources (HR) documents about applicants that were not hired. In certain legislative regions, these documents must be disposed of within short time frames, for example, within a month after the final decision. Failing to do so can result in substantial fines.

Construction plans, design documents, and quality records for complex infrastructures, alternatively, have retention periods that are counted in decades rather than months. Retention periods of 30 or more years are not unusual in such contexts. Note that such complex rules are not only specific to document types or business applications, but can also be different across countries with different laws, which can make the required setup in a multinational organization even more complex.

Legal holds overlay the general rules for retention of different types of information based on type, owning organization, and so on. The legal holds apply whenever business operations documents are subject to litigation, and are put under a hold order issued by a court of law to prevent their disposal before the end of the investigation.

IBM Information Lifecycle Governance (ILG) solutions help change the *keep everything* approach by delivering capabilities that the various information stakeholders, including legal departments, IT, records management, privacy officers, and business users, can use to address information governance issues in their domain. These IBM solutions help to drive down the costs and risks associated with over-retention by using tools that enable the following outcomes:

- IT staff are enabled to understand and manage, by system and employee, the content collection, archiving, and retention criteria, and the procedures established by the organization. Furthermore, they can implement an archiving program that reduces cost by reliably retaining what is important and required, and deleting what is not.
- Attorneys and paralegals are enabled to automate legal hold processes, and coordinate evidence collections across the organization to respond to requests for information more quickly and cost effectively.
- Records managers are enabled to develop and manage global retention policies, and to coordinate compliance and disposition across multiple systems and jurisdictions.

- Privacy officers are enabled to assess and communicate privacy duty by data subject and data location, including overlapping obligations.
- IT staff is enabled to determine which systems appear to have the highest cost and risk profiles, and to enable them to address management of systems by information value.

These capabilities generate a more complete picture of the information inside an organization, and enable information management decisions based on fact and certainty. With this confidence comes the ability to implement a defensible disposal program that can have a real effect on the amount of data under management and the associated cost and risk.

The IBM Value-Based Archiving solution, and in particular the IBM Content Collector family of products, support defensible disposal efforts with a larger set of capabilities that attack the data growth problem at the source system. This approach immediately reduces the amount of information debris under management inside an organization, and the cost and risk associated with over-retention.

8.1.2 Information lifecycle governance applied to SAP systems

The previous sections of this chapter explained the business requirements that justify the need for an organization-wide content management strategy that helps to meet the challenges imposed by data growth, discovery needs, and rules for legal compliance when handling business structured and unstructured information.

As mentioned previously, the information contained in an SAP system, which in the majority of cases is not just one source but a *key* source of business operations information, must not be treated as an isolated silo. It has to be integrated into the overall strategy.

Structured information in SAP systems

Managing and controlling the size of SAP applications provides sustainability of system performance while monitoring infrastructure cost.

Application size management is primarily performed by moving infrequently used or obsolete business transaction data, which is no longer involved in active business transactions, to less costly storage while maintaining transparent access to such data. Examples include data only needed for summary reporting or infrequent auditing purposes.

The SAP system provides the tools to archive data that is declared *business complete*, and store these in a proprietary aggregation format in archive files in a secure ECM repository. The SAP system also provides transparent access to the archived data directly from the ECM repository. The effects of this maintenance operation are two-fold:

- System performance is kept under control by limiting table sizes within manageable boundaries.
- Storage cost on high-performance storage that it is needed for immediate transactional processing of table data can be reduced, because infrequently used data is moved to less-costly storage.

The management of SAP data too often exists as an IT island, disconnected from the rest of the information management of the organization, retention, and information governance standards. However, as SAP implementations grow, and as new SAP enterprise resource planning (ERP), customer relationship management (CRM), and other component modules interact with more of the overall organization, an SAP archiving solution needs to be more comprehensive and interconnected.

This approach helps ensure proper information governance for SAP system data and content, both within and outside of the SAP system. In fact, the need for scheduled, defensible disposition of SAP data and content often drives the need for archiving itself. This motivation is on a par these days with infrastructure cost reduction and overall business efficiency drivers. For more information, see 8.4, “Data governance: Managing growth and compliance” on page 221.

Unstructured information in SAP systems

Unstructured information is associated with most SAP business processes:

- Invoices in an accounts payable process
- Purchase orders and billing documents in a supply chain or financial application
- Generated reports summarizing business transactions in just about any application

This information frequently needs to be accessed from outside the SAP application:

- Customer service department employees need to see all of the information about the previous transactions of a customer. They want to provide the customer with a digital copy of their bill to reduce mailing cost.
- As part of a litigation, the legal department, which has no direct access to the SAP system, has to provide all of the contract documents from a given time period, matching certain textual search criteria, to fulfill a court order.

The critical business value in these examples is the ability to provide *all* business information, both structured and unstructured content, to the correct person at the correct time. Providing access to unstructured content directly within the business process enables the person to make the correct decision quickly. For details about some solution use cases, see 8.2, “ECM for SAP use cases and solution architecture” on page 196.

In summary, both the structured and the unstructured information in an SAP system are an integral part of the overall information landscape of an organization, and therefore must be managed in the same infrastructure of ECM, following the same strategies, and the same governance rules as non-SAP related content.

8.1.3 IBM proposes a base structure of an integrated ECM solution

The previous sections of this chapter describe how an organization-wide content management strategy, including archiving and defensible disposal of information, extends the value of an organization’s SAP landscape. Several examples demonstrate how to enhance the implementation of core end-to-end business process solutions by seamlessly integrating SAP system components with IBM ECM components.

The following list summarizes these value points:

- The IBM ECM/ILG solution extends the SAP business solution by providing the correct information at the correct time within all business processes across an organization.
- A strictly value-based archival and disposal model enables an organization to maintain system performance and reduce infrastructure cost.
- An organization-wide retention and hold model agreed upon by all stakeholders provides a framework for maintaining full compliance with complex governing rules under a common defensible disposal strategy.
- Cross-system, end-to-end solutions that implement complete business processes can be achieved with the seamless integration of SAP business modules and IBM ECM products, such as IBM Datacap for capture and IBM Content Navigator as a common ECM front end.

The IBM solution is integrated and certified by SAP. Moreover, IBM is rated as a market leader in this space by the leading industry analysts. For more of a focus on the SAP archiving process and its integration with IBM ECM portfolio components, see 8.2, “ECM for SAP use cases and solution architecture” on page 196. It has more details about the core use cases, and shows how organizations can use IBM Content Collector for SAP Applications to implement them.

The core use case can be extended with state-of-the-art document capture functionality implemented through IBM Datacap. Additionally, IBM Content Navigator is introduced as the document-centric integration platform of ECM.

8.2 ECM for SAP use cases and solution architecture

This section shows how IBM ECM portfolio software can enhance the value of an SAP business infrastructure with document and data archiving support. IBM Content Collector for SAP Applications plays a central role as the gateway product that operates between the SAP system and the ECM environment.

The integration with business applications, such as an accounts payable solution, is described in 8.3, “Business process enhancements through ECM for SAP solutions” on page 207.

Before describing the use cases for SAP archiving in detail in 8.2.2, “SAP archiving use cases” on page 197, the following high-level use cases need to be introduced:

- Archiving data and documents originating from SAP systems.
- Making documents already stored in an ECM system available to SAP users by linking them to SAP Business Objects.
- Enabling SAP users to access documents stored in the ECM repository outside the SAP graphical user interface (GUI). Such documents can be anything, from scanned images, faxes, forms, and email to documents originating in electronic form. The documents stored in the ECM repository by an SAP system can also be made available to other non-SAP applications and capabilities of the ECM platform, such as classification, records management, and e-discovery.

8.2.1 SAP archiving standards

Starting with SAP Release 2.2, SAP clearly recognized the need to enhance the SAP software through a robust archiving facility, both for offloading the operational SAP database and processing business documents that are in an external archiving system. With SAP Release 3.0, all key applications and modules offered archiving functionality that was included as part of the standard SAP package. The SAP Business Framework provides the interfaces necessary to integrate these external functions with SAP systems.

SAP ArchiveLink

The primary interface for integrating storage and content management systems into an SAP system is called SAP ArchiveLink. It was introduced with Release 2.2 and enhanced in subsequent releases.

Additionally, the SAP Hypertext Transfer Protocol (HTTP) Content Server interface was defined in SAP Release 4.5 as a subset of SAP ArchiveLink that focuses on content rather than storage management. HTTP Content Server interface is a general, cross-application

interface that connects the enabled SAP business applications to a content server, and enables them to process documents in logical content repositories.

This content server can be a database, a file server, an SAP system, or an external archive. The following list describes the supported SAP ArchiveLink 4.5 functions:

- HTTP Content Server interface
- Business Application Programming Interface (BAPI) for bar code-based archiving scenarios and creation of SAP work items
- Object Linking and Embedding (OLE) functionality for storing inbound documents or PC files, and starting external viewing applications on Microsoft front ends

SAP ArchiveLink has not been changed since SAP Release 4.6. Note, however, that the current SAP ArchiveLink specification has made the requirement to support OLE optional, so vendors might drop support for OLE in future versions of their SAP ArchiveLink software.

Although SAP ArchiveLink is focused on the management of content and storage, it is not suited by itself to address compliance use cases, such as decommissioning existing systems, managing data retention rules, and collecting and preserving data for legal cases through legal holds.

SAP Information Lifecycle Management

SAP introduced SAP Information Lifecycle Management (ILM) to manage the lifecycle of productive or archived data and documents. With ILM, data and documents are stored in an ILM Web Distributed Authoring and Versioning (WebDAV) Server, and the following retention management capabilities are provided:

- Placing retention policies on documents and data to determine how long they need to be kept in the archive to comply with governmental rules and regulations
- Placing legal holds on data and documents that are needed in the context of a legal case
- Enabling controlled disposal of documents that have reached the end of their retention period and are no longer needed as part of legal holds

For more details about ILM, see 8.4, “Data governance: Managing growth and compliance” on page 221.

8.2.2 SAP archiving use cases

This section considers the various use cases for SAP archiving along several important dimensions:

- Data versus document archiving
- Early versus late and simultaneous archiving
- Use of bar code technology
- Access to documents from outside of an SAP system

Data archiving

Business records that are no longer needed on a daily basis can be packaged in an SAP-defined format called Application Development Kit (ADK) file, and archived using the SAP ArchiveLink or HTTP Content Server protocols. By doing so, organizations can keep the SAP database at a manageable size.

This archived data is still accessible by an SAP system in a transparent manner, while at the same time reducing storage costs, increasing productivity, and improving system performance. For more details, see 8.4, “Data governance: Managing growth and compliance” on page 221.

Document archiving

For legal and internal policy reasons, companies must keep documents pertaining to their business operations for a certain period of time. Filing the documents in paper form has several disadvantages because of their physical nature. For example, they must be duplicated when more than one user needs them. Tracking their physical location reliably can be cumbersome and error-prone.

Processing documents electronically has the following benefits:

- Organizations can use cost-efficient storage media.
- All authorized users can access the documents without being delayed by conventional archive inquiries.
- Several users can access the same documents at the same time.
- Disaster recovery (DR) procedures can be fully automated. Workflow processes, such as an approval procedure, can be defined consistently organization-wide, and can be fully automated.

The following types of documents can be identified:

- Incoming documents. These documents include, for example, supplier invoices that reach the company by way of mail or telefax, and are typically stored as digitized images.
- Outgoing documents. Documents created electronically, printed, and sent to their respective recipients. Archiving the electronic form of these documents enables for fast retrievals for customer inquiries or audits.
- Reports and print lists. A specific type of outgoing documents, generated by an SAP application, that are usually printed. The use of an archive makes a physical printout obsolete. The electronic journal, as opposed to its paper equivalent, is easily searchable, and can even contain hyperlinks to other documents to enable convenient cross-reference.

  For example, a specific entry in a document journal might refer to a scanned original document. Furthermore, the archived lists can be used as input to other applications, such as data mining systems, for advanced analytics purposes.
- Desktop files. Documents that are created by applications, such as office applications or common agent services (CAS) applications. They can be archived and later accessed by an SAP system using Desktop Office Integration (DOI), based on SAP Business Connector to SAP Central Instance (BC-CI) integration, or using the document management system (DMS).

Early archiving

In early archiving, an incoming document is first captured into a repository, for example, IBM FileNet Content Manager. It is then made available for processing by SAP applications before the associated SAP business object, for example, a sales order, is created. The incoming document drives the creation of the SAP business object. This approach eliminates paper processing from the beginning of the process, and separates the scanning process from the creation of the business object.

When the document is captured and assigned to an SAP document type, an SAP user is notified that a new document is ready to be processed. The user then creates the business object associated with the document type. The notification process uses SAP Business Workflow, and a link between the business object and the document in the repository is established.

Simultaneous archiving

In simultaneous archiving, all document entry and SAP object processing steps are carried out by the same SAP user. Overall, the process is the same as in early archiving except that the SAP work item, which is created at link time, is assigned to the current user.

Late archiving

In late archiving, the creation of the SAP business object comes first, and linking to the corresponding incoming document happens later in the process. In practical terms, this process is like a traditional paper-based process. The SAP business object already exists when the incoming document is captured into a repository, and a link between the business object and the incoming document is established.

Use of bar code technology

Document archiving can include bar code linking, whereby a bar code that is applied to the scanned document is captured by an SAP ArchiveLink-compliant repository and uploaded to the open bar code table of the SAP system for linking. Bar code scenarios are particularly suited when existing paper-based processes should not be disturbed by the introduction of scanning and archiving technology. Documents are still distributed in paper form to the point where the business object is created and the linking occurs.

Similar scenarios where the electronic processing and the physical handling of paper remain uncoupled for some time include the handling of receipts in expense processing, and the processing of contracts that require physical signatures at completion time. In both cases, the use of bar codes helps to bring paper and electronic records together.

Access to documents from outside an SAP system

The archiving use cases described so far are within the context of the SAP ArchiveLink protocol and specification. One use case exists, however, that goes beyond what is offered by SAP ArchiveLink.

All SAP ArchiveLink operations are based on the concept that the SAP database represents a master index catalog of all documents, including those whose content resides in the external archive. SAP ArchiveLink therefore passes no business information to the external archiving system except a Universally Unique Identifier (UUID), which the SAP software uses to identify a document throughout its lifecycle.

For many SAP users however, it is a requirement to access archived documents, including those from non-SAP contexts, independently of the SAP system (without using the SAP GUI). To support federated searching of the external archive by business information, such as customer number or fiscal year, the corresponding document attributes must be transferred from the SAP system to the external archive to be searchable there. This process is commonly referred to as *index transfer*.

This transfer of business information from SAP Business Objects referring to the corresponding documents in the external archive, to the external archive as document attributes, is usually achieved by proprietary function modules. These modules must be installed on the SAP system and configured to collect and transfer the required document attributes from the SAP system to the archive.

Implementing use cases with IBM ECM products

The previous sections describe the use cases relevant for SAP archiving, some of which even go beyond what is defined by the SAP ArchiveLink standard. The following IBM ECM portfolio products can be used to implement these use cases:

- IBM Content Collector for SAP Applications
- IBM Datacap
- IBM Content Navigator

8.2.3 Architecture of IBM Content Collector for SAP Applications

The architecture of IBM Content Collector for SAP Applications enhances the SAP business infrastructure with data and document archiving. With the components in this architecture, you can archive and view data and documents, and print lists. You can also link archived documents with SAP business objects.

Figure 8-1 shows the high-level architecture of IBM Content Collector for SAP Applications. Content Collector for SAP acts as a gateway that interfaces between the SAP system and the ECM repositories.

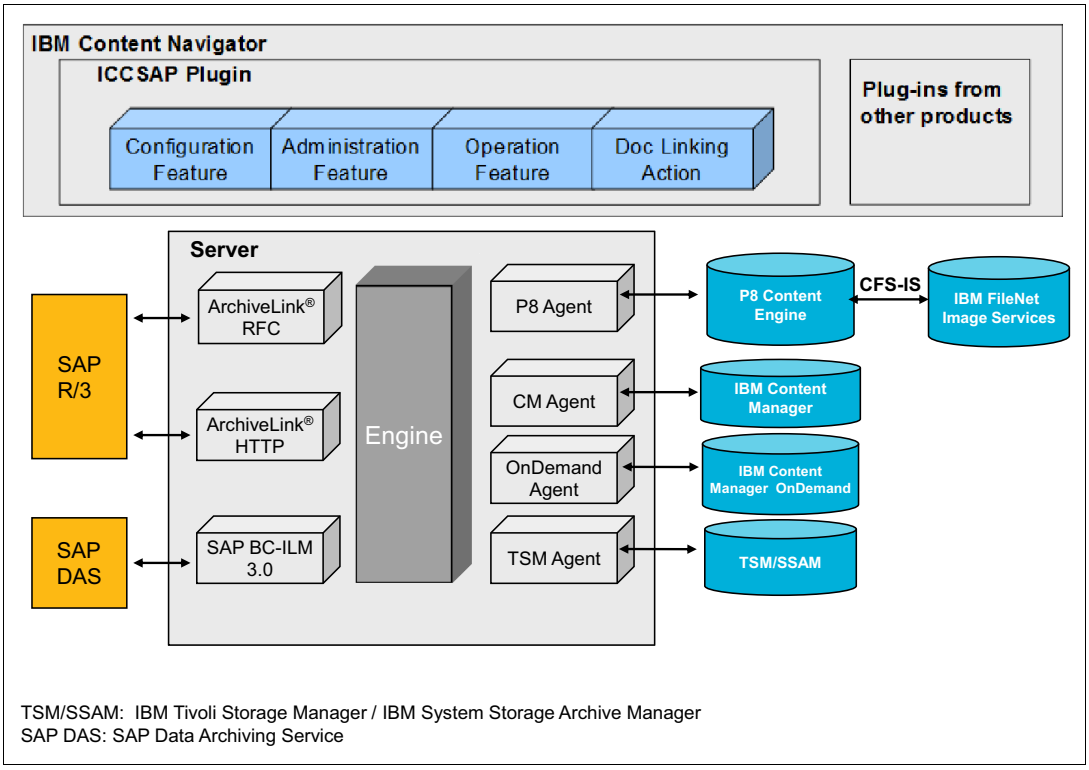

Figure 8-1 IBM Content Collector for SAP Applications product components

IBM Content Collector for SAP Applications supports the archiving of data and documents using both versions of SAP ArchiveLink into four back-end systems:

- IBM FileNet Content Manager
- IBM Content Manager Enterprise Edition
- IBM Content Manager OnDemand
- IBM Tivoli Storage Manager

IBM FileNet Image Services can be used as a back-end for IBM Content Collector for SAP Applications, accessed transparently through the Content Federation Services agent for FileNet Content Manager.

The centerpiece of the architecture is the engine of the IBM Content Collector server, which distributes the incoming requests from an SAP system to the back-end systems and returns the responses back to the requesting SAP system through the use of dispatchers and agents.

The user interface for IBM Content Collector for SAP Applications is based on IBM Content Navigator, the common user interface of all IBM ECM repositories.

The following paragraphs provide more detail about the components of IBM Content Collector for SAP Applications.

Server

The server of IBM Content Collector for SAP Applications is the central component that handles all operations. It contains the engine and the components that connect to SAP and to all back-end systems. This set of components, operating in its own port range, is called an *instance*.

The engine processes all inbound and outbound communication in a bidirectional way. Communications include archival and retrieval requests from SAP and their translation into search, retrieval, and archival requests to the attached ECM repository also.

The connections to SAP on the left side of Figure 8-1 on page 200 are implemented as dispatchers, which can be started in a configurable number to address different workloads. The SAP ArchiveLink based dispatchers are RFC and HTTP based, depending on whether the newer 4.5 version or the previous 3.1 version of the protocol is used. The SAP ILM dispatcher is also HTTP based, but uses the WebDAV protocol.

The following archiving protocols that use SAP Business Connector are supported:

- SAP ArchiveLink (BC-AL) versions 3.1 and 4.5
- SAP HTTP Content Server (BC-HCS), which is a subset of the full SAP ArchiveLink specification, comprising the pure server-to-server communication without any front-end scenarios.
- SAP Information Lifecycle Management (BC-ILM), which is an SAP extension of the WebDAV protocol. WebDAV is the protocol for the Extensible Markup Language (XML) Data Archiving Service, referred to as SAP DAS in Figure 8-1 on page 200.

The connections to the repositories translate the repository-agnostic requests from the engine into repository-specific requests by using the respective API. For each back end, a separate agent exists. The number of agents is also configurable in order to adapt to different workloads.

The IBM Content Collector Server is also scalable horizontally by starting additional instances of the entire collector server, where each instance can operate independently from the other by assigning individual port ranges. This mode of operation is also used when multiple SAP systems use the archival services that are provided by IBM Content Collector for SAP Applications.

Components related to the client interface

IBM Content Collector for SAP Applications adopted IBM Content Navigator as the basis for all of its client functionality as described in this section.

For more information about IBM Content Navigator, its integration and customization capabilities, and its extensibility, see *Customizing and Extending IBM Content Navigator*, SG24-8055.

As shown in Figure 8-1 on page 200, IBM Content Collector for SAP Applications provides a single client interface, which offers functionality to perform the following actions:

- Configure an IBM Content Collector instance.
- Create and administer profiles for archiving, document linking and index transfer.
- Administer, schedule, execute, and monitor tasks based on those profiles.

It also integrates with the advanced document viewing capabilities of IBM Content Navigator, based on IBM Daeja™ ViewONE technology.

The plug-in architecture of IBM Content Navigator allows for a seamless integration with IBM Content Navigator components of other ECM products such as IBM Case Manager or IBM Datacap.

IBM Content Collector for SAP Applications uses the IBM Content Navigator plug-in architecture to provide the functionality mentioned previously. See Figure 8-2.

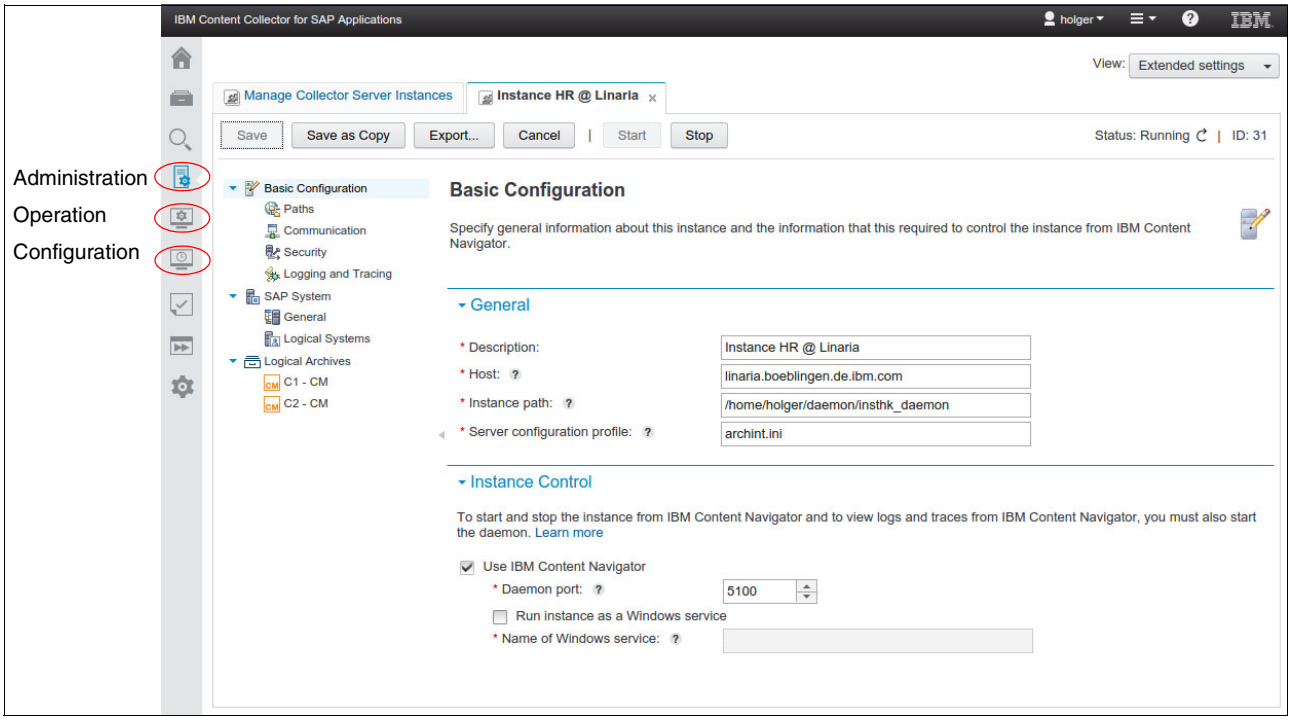

Figure 8-2 IBM Content Collector for SAP Applications configuration using IBM Content Navigator

Configuration feature

In the instance configuration feature, administrators create and maintain all IBM Content Collector for SAP Applications instances. One configuration feature (in one Content Navigator instance) can maintain instances on multiple hosts, covering multiple SAP systems (in general one per instance), and multiple back-end repositories.

One instance configuration collects all the necessary information about the instance itself and the characteristics of the participating systems on the SAP and the back-end repositories side such as these examples:

- User credentials for the SAP system and the back-end repository
- Communication details:
  - Communication protocols and ports
  - Security options and certificate handling
- Logical archive configuration:
  - Document classes that correspond to SAP document types
  - Folder structure of the logical archive
  - Document metadata mapping
  - Viewing options based on document type and MIME type information
  - Repository-specific configuration items

The complexity of the logical archive configuration depends on the choice of the back-end repository and, therefore, the functional capabilities that are supported by that repository.

The configuration feature uses active connections to the SAP system and to the back-end repository to provide as much information about the running systems as possible. With this information, you can perform consistency checks across the entire configuration as early as possible, and the information significantly reduces the number of configuration errors that might occur in a purely manual operation.

Document viewing options

Depending on the SAP business application that is used, the types and formats of attached documents can vary greatly. With that, the need for a highly configurable set of options for viewing these documents is necessary. Traditionally, the SAP system includes many options for viewing documents from within the SAP GUI, but external viewers can also be added and integrated.

The availability of IBM Content Navigator as the common user interface for the content repositories adds even more flexibility to this use case. IBM Content Navigator includes the IBM Daeja ViewONE viewer, which is highly configurable for many viewing scenarios. The logical archive configuration section in the configuration feature provides the options to activate these viewing methods.

Administration feature

In the IBM Content Collector for SAP Applications administration feature, for each SAP archiving use case (described in 8.2.2, “SAP archiving use cases” on page 197) profiles are created that describe the operation performed and the content parameters of that operation, including the following information and details:

- Information about the source of the documents that are processed, that is, whether they are from an external source, such as a capture solution destination, or whether they already reside in a content repository.
- Information about specific document linking methods, to distinguish between these
  - SAP barcode processing
  - The creation of SAP work items that trigger business specific SAP workflows
- Details about the mapping of attributes that synchronize document properties with corresponding business object metadata in the SAP system

Under this general setup, three types of profiles are configurable:

- Archiving profiles
- Document linking profiles (which also cover the special case of manual linking)
- Index transfer profiles

Archiving profile

An archiving profile is used to describe the necessary transactions to transport a set of documents from an external source into the repository, and simultaneously create the association of these documents with corresponding SAP business objects. (In that sense, every “archiving” operation is implicitly combined with a “linking” operation.)

The archiving profile must provide the following set of information:

- The origin of the documents

  These can be from an external capture solution, or from some other type of application that deposits documents into a file system location that the archiving process monitors for new entries, which are then archived.

- The target document class on the back-end repository

  This controls the set of metadata that will be assigned to the document.

- The linking method and the corresponding document type on the SAP system side, if the linking method is Create Work Item.

Two basic linking operations are available, as specified by the SAP ArchiveLink protocol:

- Creating an SAP workflow work item

  Based on content that is placed onto a work queue in FileNet Content Manager or into a work-basket in IBM Content Manager Enterprise Edition repositories, SAP workflow work items that are based on standard work tasks can be created and be configured to start the appropriate SAP transaction according to each SAP document type.

- Creating SAP external bar code entries

  Based on barcode information that is present as document metadata, a link to the associated business object is established if the values of the open external and the internal barcode tables match.

Document linking profile

A linking profile contains all the necessary information for performing the operation of associating documents that are already present in the document repository with SAP business objects in the SAP system.

Similar to the archiving profile, the linking profile also distinguishes between the two main methods of creating the SAP-based information:

- Create a work item, as described for the archiving profile (in “Archiving profile”).

  In addition to the triggering of a specific workflow on the SAP system side, additional metadata (taken from the document’s metadata) can be specified. These are then added as metadata to the SAP business object during the document linking action.

- Create external SAP bar codes (as described in “Archiving profile”).

Index transfer profile

In some business scenarios, you might need to access documents that are related to SAP Business Objects outside of the SAP system, but within the business context. Documents that are created in an SAP system, and archived in IBM Content Manager Enterprise Edition, IBM Content Manager OnDemand, or FileNet Content Manager by using IBM Content Collector

for SAP Applications, do not contain any searchable business data-related attributes by default. To support such a scenario, IBM Content Collector for SAP Applications provides a function to transfer business data from SAP to the ECM repository as document properties. This capability is called index transfer.

Through the use of folders, a structured document hierarchy that, as an example, exists in a contract management solution in SAP, can be mirrored in the content repository, providing seamless access to documents from within and from outside of the SAP system.

In an index transfer profile, you specify the required parameters to synchronize selected metadata between business objects in the SAP system and their associated documents in the content repository (for repositories that do support enrichment with metadata).

An index transfer profile includes the following information:

- The SAP business objects tables involved (for example, a BKPF table in an FI application). These can also be user-defined tables.
- The document types that should receive the metadata.
- Mapping information that associates SAP business object metadata with configured properties of the document types in the repository. The corresponding data types must be compatible and commensurate.

Operation feature

The operation feature of the IBM Content Collector for SAP Applications plug-in is the place where the profiles that are created in the administration feature are put to use in day-to-day operations.

Based on these profiles, tasks are created that describe the operational aspects:

- Task schedules

  Each task can be configured to run only once or repeatedly.

  Recurring tasks are assigned a task schedule, describing on which day of the week at which hour and minute the task is scheduled to run. With this function, you can plan task resource usage according to anticipated system load and resource availability.

- Task status

  Planned tasks can be monitored for their individual progress, and for their overall status at task termination, by indicating the number of processed items and potential error status.

  Recurring tasks can be suspended and resumed.

- Task auditing

  For all tasks that are administered through the operations feature, the task manager component of IBM Content Navigator provides auditing facilities to document the task activities.

Client API

External applications, for example document capture solutions such as IBM Datacap, communicate with the Collector Server by using a public client API. External applications can integrate with IBM Content Collector for SAP Applications through the use of this client API. The client API supports the archival of documents and the subsequent action of either creating SAP Work Items or sending bar codes to link documents to SAP Business Objects.

IBM Datacap is integrated in this way, and integrations with IBM Business Partner applications were also implemented. For more details, see 8.2.4, “IBM Datacap” on page 206.

Figure 8-3 on page 207 shows the IBM Datacap integration schematically. A document is first scanned and processed by IBM Datacap, which extracts certain information, for example, an invoice number, from the scanned document. This information is written to a properties file, and the custom IBM Datacap action calls the client API and passes the relevant information to it. Metadata, such as the invoice number, can be stored in the archive or the SAP system along with the scanned document.

Scalability and high availability

IBM Content Collector for SAP Applications architecture is designed for scalability and high availability (HA).

The IBM white paper *IBM Content Collector for SAP Applications: Sizing, Configuration, and High Availability* explains how to size an IBM Content Collector for SAP Applications solution, and how to scale it and configure it for HA. It is available at the following web page:

`http://www.ibm.com/support/docview.wss?uid=swg27036773`

The SAP-driven scenarios for outgoing documents can be made highly available by using standard techniques, such as a load-balancer for the HTTP traffic.

8.2.4 IBM Datacap

The ability to efficiently convert paper records into digital form by scanning and making them available to SAP systems as an incoming document is a key ingredient in any initiative for reducing manual paper processing.

IBM Datacap quickly and easily captures, manages, and integrates enterprise content while extracting critical information. IBM Datacap offerings include easy-to-use customization with high-volume document capture. The IBM Datacap product family consists of many components. In the context of this chapter, consider the following components:

- IBM Datacap is based on a service-oriented architecture (SOA) capture and automation solution that includes both web and thick clients. It is a document ingestion product.
- IBM Datacap Studio is used to configure the IBM Datacap solution by defining and assembling the document hierarchies, recognition zones and fields, rules, and actions.

Like IBM Content Collector for SAP Applications, the user interface of IBM Datacap is based on IBM Content Navigator.

For more information about IBM Datacap, see *Implementing Document Imaging and Capture Solutions with IBM Datacap*, SG24-7969-01.

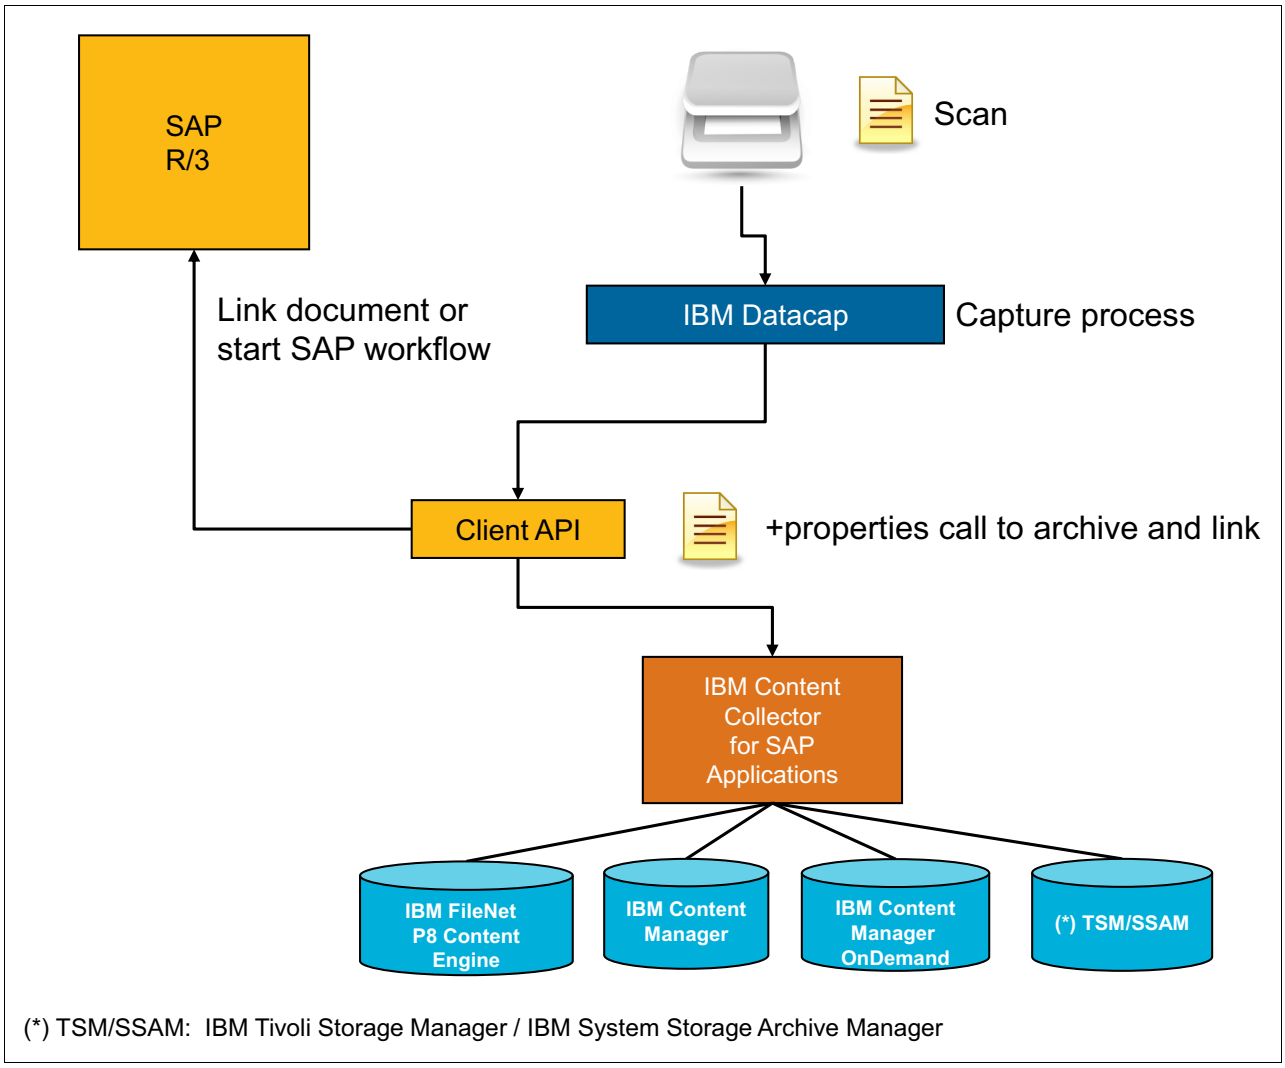

Figure 8-3 Capture and archive process integrating IBM Datacap and IBM Content Collector for SAP Applications

8.3 Business process enhancements through ECM for SAP solutions

Even in a “paperless world”, the majority of today’s business processes are document-centric. A tight and smart combination of the strength of SAP in ERP systems with the sophisticated capabilities of the IBM ECM infrastructure enables business users to benefit from an integrated cross-system workflow management. An important point to understand is that those solutions and enhanced features go significantly beyond the SAP ArchiveLink specification, which only requires a simple imaging and archiving system.

8.3.1 Objectives of a document-oriented workflow management

The operation of a document-oriented workflow management system usually pursues one or many of the following objectives:

- Enable end-to-end processing
- Increase the degree of automation

- Decrease error rates
- Reduce handling expenses
- Relieve users from routine tasks
- Adhere to compliance rules

Examples include HR file management, invoice processing, contract management, processing of order sheets, and delivery receipts.

8.3.2 SAP-centric versus ECM-centric process management

Basically, two fundamental approaches are available for an integrated workflow management:

- ECM-centric. A separate business application that communicates through interfaces with the SAP system. This application is typically an integral part of the ECM system, either a workflow component or a case management application.
- SAP-centric. The business application is deployed directly on the SAP system.

The ECM-centric process management configuration is usually covered by functions of the ECM systems. Their archiving and imaging capabilities are needed to satisfy the SAP ArchiveLink use case, but they are typically full-featured ECM systems. Configuration and maintenance are in general easier than they are for the complex SAP ERP system. However, the main disadvantage is that SAP is still the leading system from a business perspective, because all of the business data is kept in the ERP system.

More than that, all related data is stored there. For instance, the processing of invoices requires access to purchase order or even vendor contract data. An accounts payable solution running inside the ECM system has to synchronize all the basic and related data with the ERP system, which process can be error-prone and difficult to manage.

Business applications directly deployed on the SAP system ensure an integrated workflow management by default. Typically, they are installed as an ABAP module within its own name space. Therefore, they behave like any other standard SAP module. The following list describes the benefits of this approach:

- The entire workflow management happens within a single system. The user does not need to learn different UIs.
- Direct access to data of the underlying SAP system exists. Therefore, no configuration or maintenance of external interfaces is necessary.
- Any relevant business data from any other SAP module required for processing an incident is directly accessible.
- Even non-SAP GUI users can easily be included through web UIs.
- Not a standard capability, but SAP Online Service System (OSS) integration can be highly advantageous.

  The SAP integrated business application should be integrated into the SAP OSS. By doing so, the support process for the business application is the same as for the SAP system itself. Ideally, a support case is directly routed from SAP level 2 support to the application provider.

Further considerations in this section are based on the SAP-centric process management approach because, based on practical experience, this is the preferred approach for SAP-centric organizations.

However, a fundamental requirement for both approaches is that the ECM system is well integrated with the SAP system. This consideration goes far beyond support for the SAP ArchiveLink interface. It includes synchronization of documents, their metadata, and embedding them into the business context between both systems.

8.3.3 Components of an ECM for SAP Solution

Solutions for an integrated workflow management typically consist of the following components (see Figure 8-4):

1. Capturing the documents
2. Processing document content and integrating it into the SAP business process
3. Filing documents and archiving them into the repository

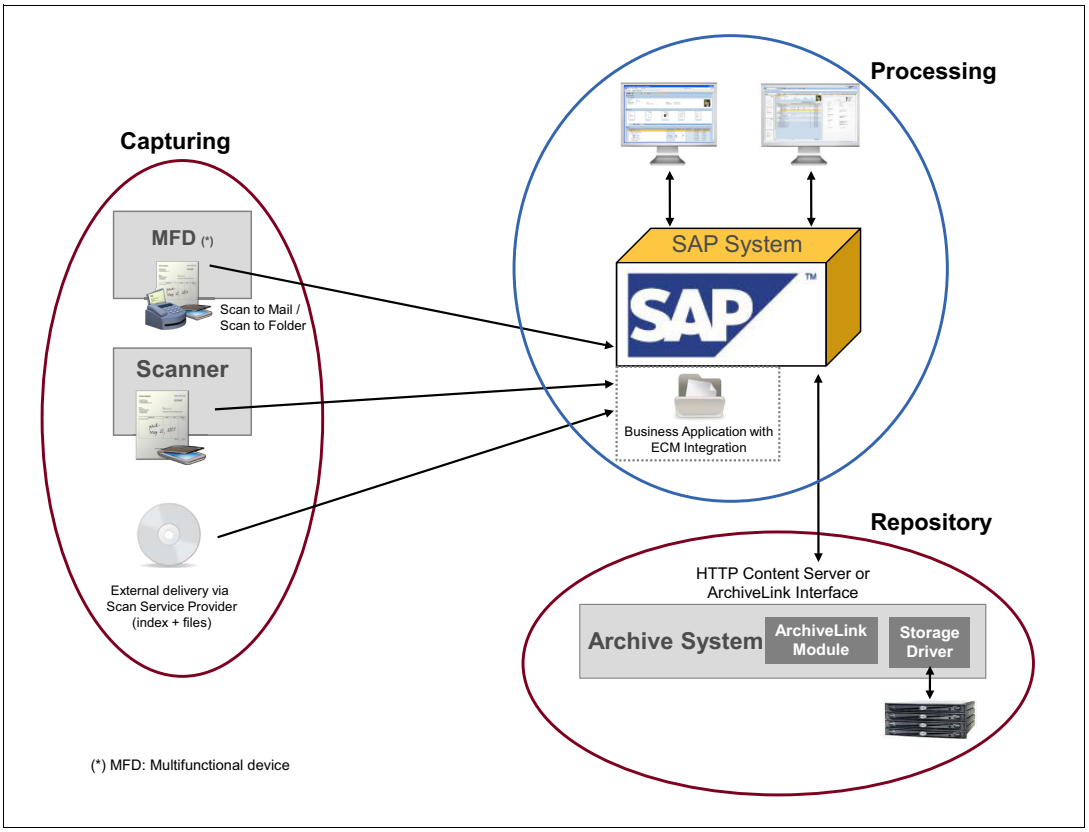

Figure 8-4 Generic components of an integrated process management

These components might be delivered either from multiple vendors or from one source. A single vendor is no guarantee of a high-quality integration, because solutions from a single vendor are often compiled through acquisitions. The quality of the integration, which is an important criteria for a seamless end-to-end process, is sometimes revealed during its usage in the field. The same concept applies to solutions where multiple vendors work together and one of them appears as the main contractor.

Fortunately, ECM systems adhering to the established SAP ArchiveLink standard enable a flexible and independent combination of capturing and repository system.

8.3.4 Capturing solution components

This section covers in more detail the individual components of the capturing solution.

Capturing

The capturing process extracts the business-relevant data from documents that are from various sources and converts them into an electronic format (see the example in Figure 8-5).

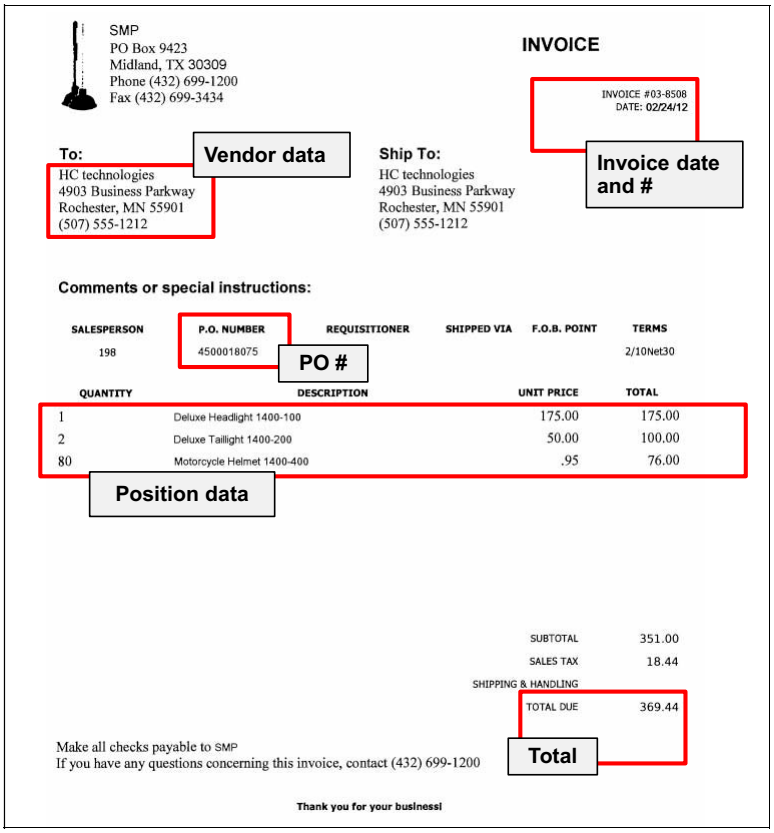

Figure 8-5 Examples of invoice data to be recognized and captured

In the example shown in Figure 8-5, the areas enclosed in red frames represent the data that has to be extracted:

- The vendor name and their postal address
- The date and serial number of the invoice
- Identification of the purchase order corresponding to the invoice
- Position data and the total sum of the invoice

The documents can be submitted either in hardcopy form, for example, invoices, delivery notes, applications, and so on, or in an electronic format, for example, through electronic data interchange (EDI). For EDI, the format conversion can be limited to a pre-processing to make it workable for further processing in the SAP system. This conversion can even be omitted if both business partners have already aligned their exchange formats. The format is standardized, for example, in the automotive industry.

Even in times of electronic documents, the capturing of business data from paper documents remains an indispensable function for the foreseeable future. For invoices and delivery notes in particular, the flawless capturing of the header and position data is a basic requirement to achieve a high throughput without manual interaction.

The recognition of the semantic content (in which position on the various forms can that specific invoice information be found) is a special requirement in a mass capture scenario. For instance, which content represents the order number, the invoice number, address data, invoice date, and so on. Of course, the position data of multi-page invoices and delivery notes has to be recognized accurately also.

In contrast, OCR recognition (capturing data and transforming it into electronic character sequences) is the easier task.

The following technologies are the basic approaches to the challenging task of recognizing semantic content:

- Free-form recognition
- Form-based recognition

Sometimes the technologies are applied in combination. such as in IBM Datacap. These technologies have a demanding theoretical background, for which a detailed description goes beyond the scope of this book.

Free-form recognition identifies the semantic meaning from the content itself. Form-based recognition concludes the semantic meaning from the position information, for example, a previously performed training mode tells the software at which location on the page to find the invoice number. The assumption here is that the vendor uses a certain standard layout for their invoices, which is typically a safe assumption.

The need for a training phase is not a challenge unique to form-based recognition. Free-form recognition requires a certain training also. The payoff for the effort spent for a precise adjusting of the recognition to a particular form is a high automation rate without any manual interaction in production. In the retail industry in particular, with its numerous vendors and a huge number of bills and receipts, a high automation rate can cause a significant increase in efficiency.

The details of mass and batch processing, such as holding or re-sorting of batches, monitoring, and so on, are not described in this chapter.

An important aspect of an integrated process management, however, is the connection with the SAP system. The transfer of the scanned documents occurs together with the extracted data in XML format, as a batch or as single documents. It can be achieved either through a shared directory in a universal and flexible way or through a direct network connection, for example, using web services.

Depending on how the invoice approval is organized, a first comparison of an invoice with the corresponding SAP purchase order data can be a reasonable approach, as shown in Figure 8-6. This approach does not substitute the final invoice processing in SAP Financials (FI). However, it can help sort out obviously incorrect documents at an early stage, in addition to correcting minor deviations, such as splitting a single order item into two invoice items.

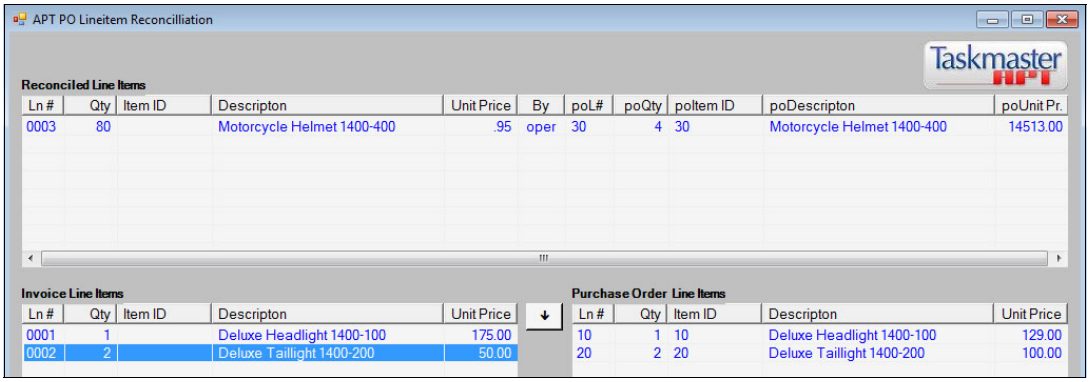

Figure 8-6 Comparison of invoice items (left side) with order items (right side) during the recognition process

Processing

The question can arise: Why introduce an additional module from a third-party vendor for accounts payable, for instance, when SAP already provides predefined workflows for invoice processing in it as standard in the SAP Financials (FI) module? Furthermore, predefined tasks already exist as part of the SAP ArchiveLink customization for the late archiving and early archiving scenarios, for example, TS3001128 and TS3001117 (see 8.2.2, “SAP archiving use cases” on page 197).

The simple answer is: The main benefit is adding the ability to extend functionality and to simplify configuration and operation. These benefits are illustrated by the examples in the following list:

- Better support for mass processing

  At first glance, the standard SAP ArchiveLink scenarios provide the capability for mass processing of inbound documents, for example, late archiving with bar code. However, by default, the SAP system is not prepared for the flexible processing of extracted data coming from a capture system.

  Posting FI records with externally captured invoice data is feasible through the SAP transaction MIRO (Enter Incoming Invoice). However, a flexibly configurable job control and queue management to buffer, control, and monitor large amounts of input data as shown in Figure 8-7 on page 213, is only available through the added third-party module.

The upper section of Figure 8-7 shows all elements of the input queue (only two are in this example) and the lower section can display the corresponding purchase order items for the selected invoice. Other options, such as displaying the invoice line items, contract data, and so on, are configurable also.

This centralized queuing and monitoring is a key capability in larger deployments. As an example, the capability to serve the invoice processing for 53 SAP systems simultaneously was the main criteria for vendor selection at a large German automotive supplier.

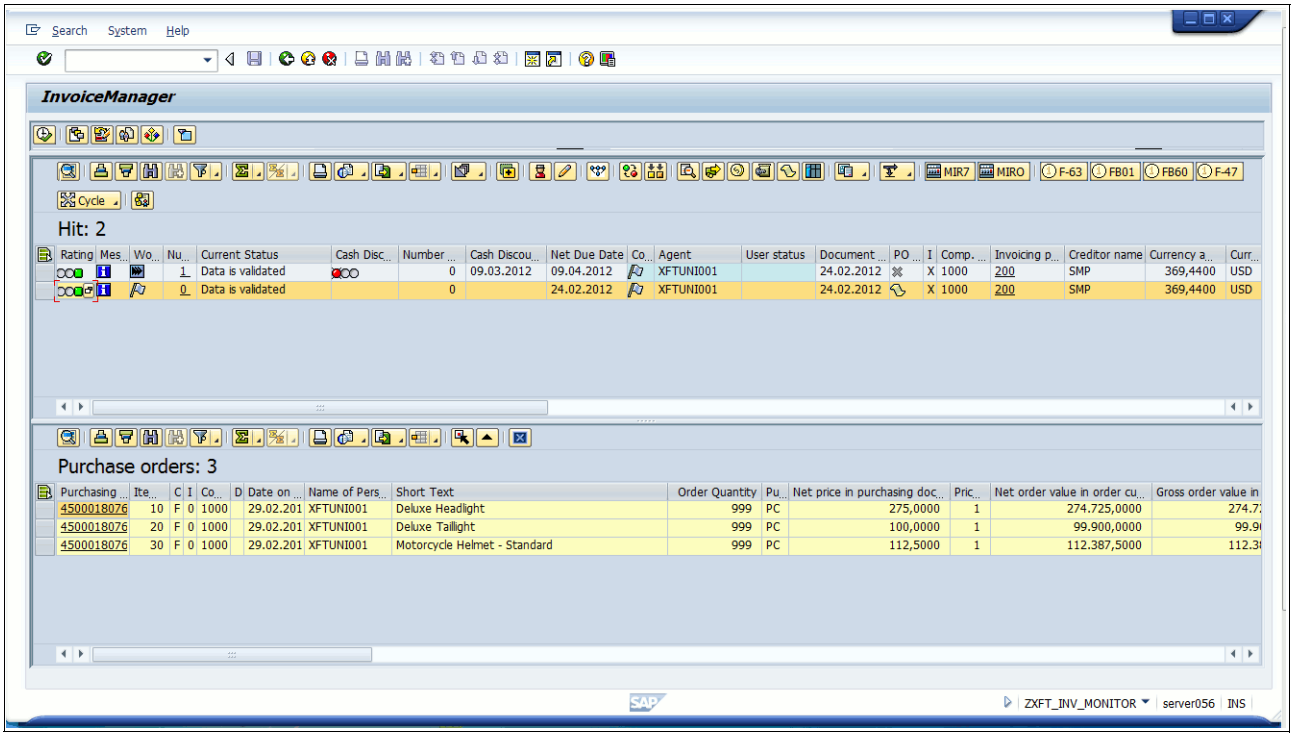

Figure 8-7 Centralized queue management and monitoring for all incoming input: Capture application example

- Simpler process configuration when compared to native SAP Workflow

  The processing of invoices requires a simple configuration of workflows in conjunction with the option to modify them easily and on short notice. For example, you can reroute the invoice processing from one operator to another, or include an additional approver, as shown in Figure 8-8.

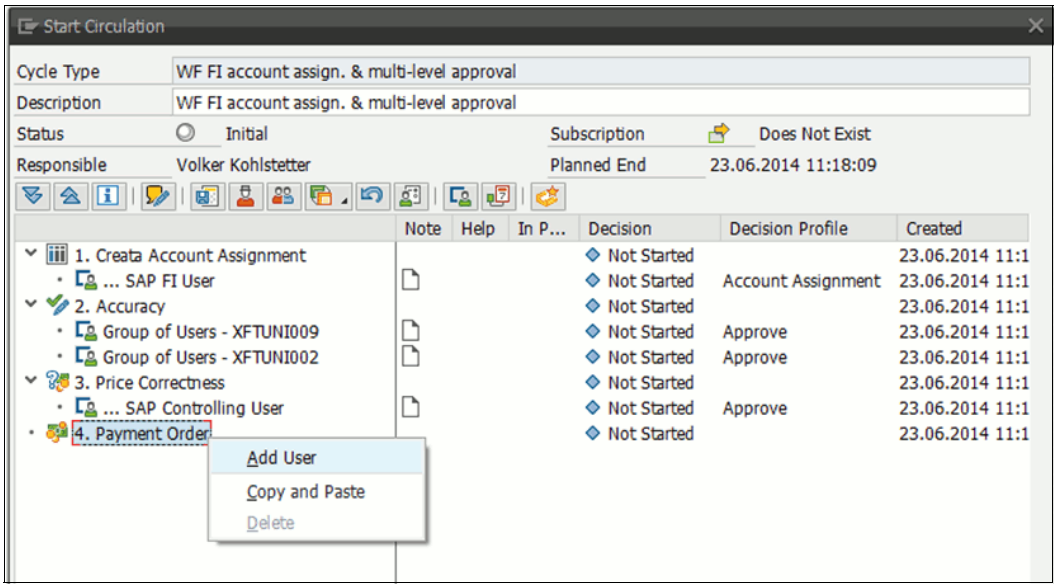

Figure 8-8 Simple and flexible configuration by adding another operator to an approval process

- Seamless integration with office systems

  Examples of this integration include dragging an email with an attachment into a personnel file to file it, as shown in Figure 8-9, or compiling response letters directly from the accounts payable solution automatically.

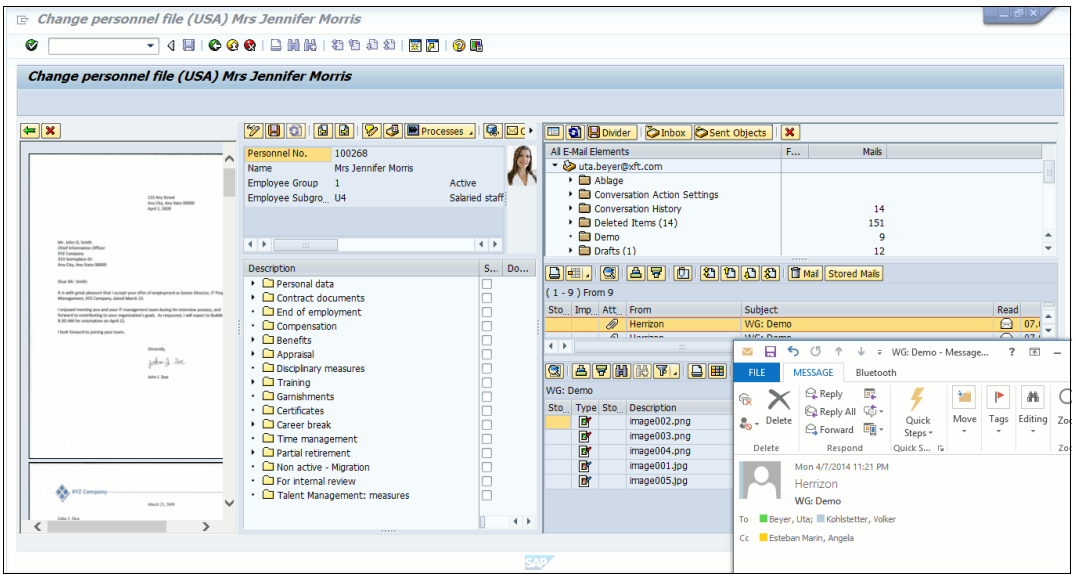

Figure 8-9 Filing an email into digital personnel file by dragging it

- Easy access to data and functions of other SAP modules directly from within the workflow application

  The resolution of any discrepancies during invoice processing often requires queries from other modules. For example, access to the vendor contract is needed to validate discount conditions on the invoice. Figure 8-10 and Figure 8-11 on page 216 illustrate this example where a vendor contract can be verified directly from within the workflow.

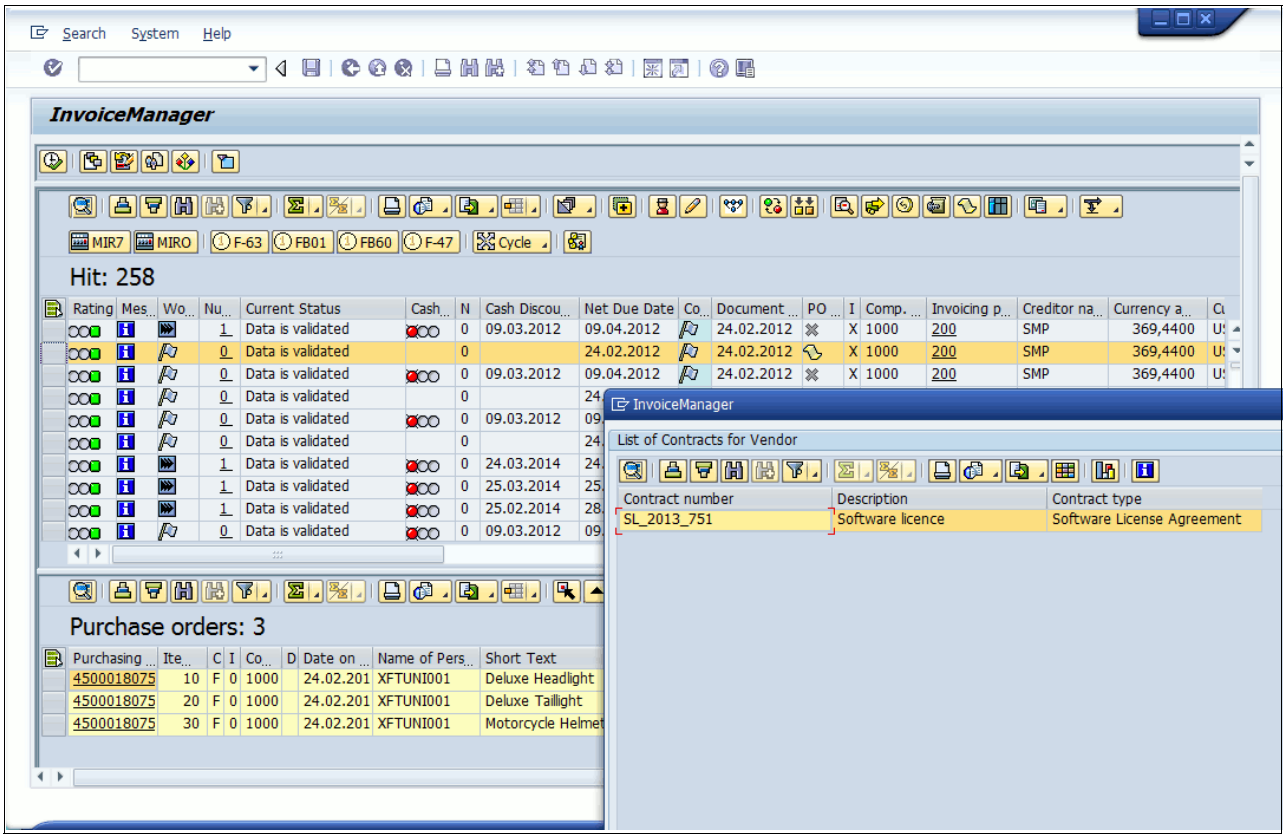

Figure 8-10 Direct access to the vendor contract from the invoice receipt list: Part 1 of 2

Figure 8-11 shows the accessed contract details.

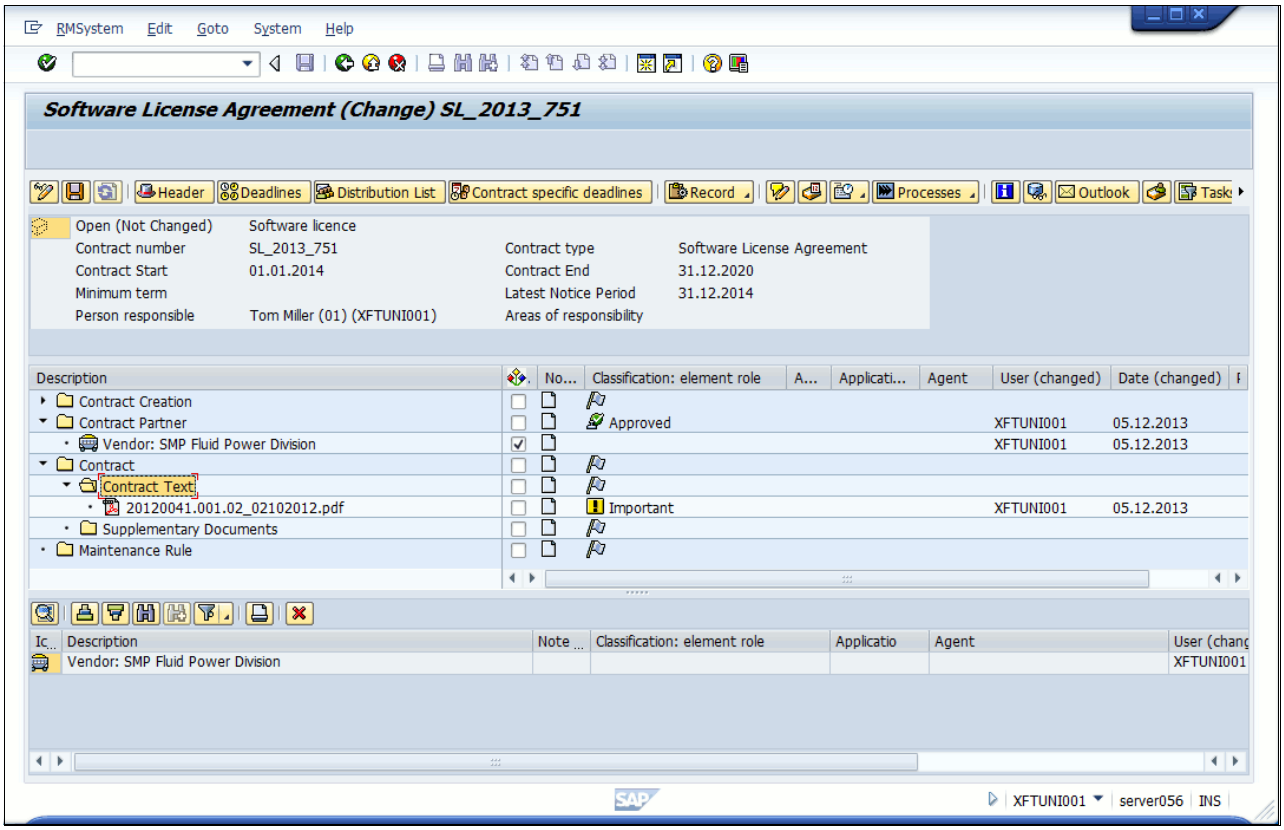

Figure 8-11 Direct access to the vendor contract from the invoice receipt list - Part 2 of 2

- Bidirectional coupling between the capture system and the repository

  A holistic and integrated process also requires a tight coupling between the capture and the archive system. Making the purchase order and vendor data available during the capture phase provides an easy way to rule out any misrouted documents in an early phase of the process, before they get into the invoice receipt book.

Repository and filing system requirements

The back-end repository and filing system must be able to fulfill a set of requirements that extend beyond the bare ability to safely store digitized content in an efficient manner. This section outlines these requirements.

Basic requirements

The main function of the repository system is to safely store the digitized documents, and to provide fast retrieval later when accessed from an SAP transaction. The technical connection occurs mainly through the SAP ArchiveLink interface. To some extent, the simplicity of this interface in its pure form limits the ECM vendors, because only low-level archiving functionality is provided.

In a pure SAP ArchiveLink solution, the only connection between the SAP system and the archiving system is the document identifier. The SAP ArchiveLink interface does not provide a way to transport any additional, potentially valuable, information, such as user information, business context, and so on. The need for extra features that add value to the solution available in today's ECM systems becomes apparent as soon as the documents filed from the SAP system are used outside of the standard use cases.

Beyond the initial low functional requirements for an SAP ArchiveLink content repository, it is advisable to take a closer look at the additional capabilities that such a repository must provide to be suited for the purpose:

- Scalability to process large volumes of documents.
- Stability of the vendor, because the purchase and operation of an archiving system are inherently designed for long-term usage.
- Flexibility connecting storage subsystems, or the capability to directly connect to a hierarchical storage management (HSM) system. IBM Content Collector for SAP Applications offers these capabilities.
- Pre-processors and user exits to enable additional operations, such as customized document format conversion. For example, you may want to convert a document from Tagged Image File Format (TIFF) to Portable Document Format (PDF).

Requirements for integrated processing

As mentioned previously, an integrated ECM SAP solution requires a direct link between the SAP system and the ECM system, which goes far beyond the capabilities defined in the SAP ArchiveLink specification. This sophisticated link basically means that the SAP business context to which the SAP-linked documents in the ECM system belong is at least partially synchronized between SAP and the ECM system.

The following examples show what this criteria means in practical terms. Consider an invoice with an attached document as a starting point (see Figure 8-12).

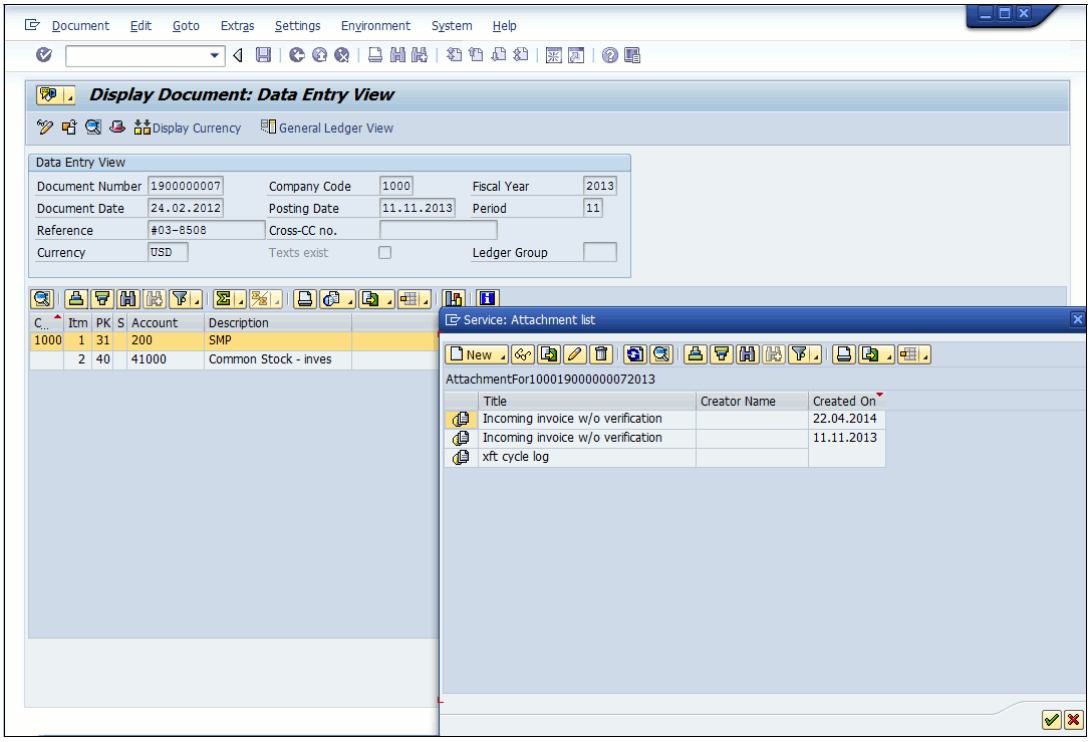

Figure 8-12 FI invoice with an attached document

In the heterogeneous SAP/non-SAP enterprise environment that is the basis of our examples throughout this book, business users work on the ECM system, where they operate on documents from the SAP system and on documents of non-SAP origin. These workers want to see not only the bare image linked to the FI invoice, but at least some of the SAP business metadata assigned as properties to the particular document.

If the workers cannot see this information, there is no chance to properly identify and retrieve the document on the ECM side, because the only default attribute is the document title (see the lower right side of Figure 8-14 on page 219).

An example of such a business context as it can be presented in an ECM environment is shown on the left in Figure 8-13. The UI is based on IBM Content Navigator. For a certain vendor, the associated invoices with their SAP document name are shown. The folder name at the lowest level reflects the SAP invoice number, and the folder contains the attached documents. The invoice folder plus a folder containing the vendor's contracts represent an excerpt of the associated business context.

Both folders are structured as sub-folders of a certain vendor. This example can easily be extended with further business information, depending on the information requirements in this context.

An important aspect to recognize is that the business information in the SAP system (the source of the information in the ECM system), and the representation of that information in the ECM system, have to be kept in sync. From a business perspective, SAP is the leading system. Therefore, all of the information accessible in the ECM system is derived from the SAP system. As information changes (some of it never, some of it frequently), it must be reflected on the ECM side in near real-time.

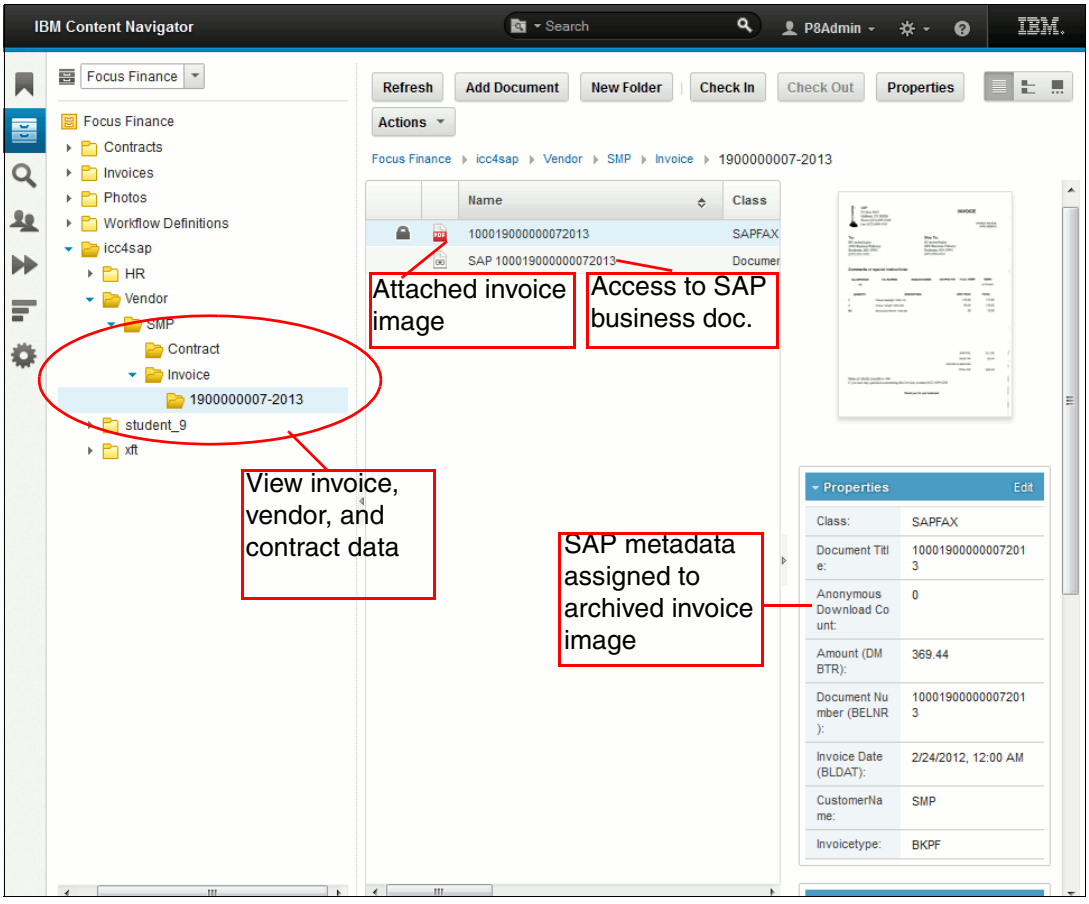

Figure 8-13 Access to SAP business data and context from within the ECM system (example)

In some cases, ECM users might need to directly access the original SAP FI business object. For this purpose, the invoice folder in our example contains a link to the FI document (see center section of Figure 8-14). When clicked, the SAP web GUI opens and the user is taken directly to the invoice record transaction without having to navigate the SAP system manually (see Figure 8-14).

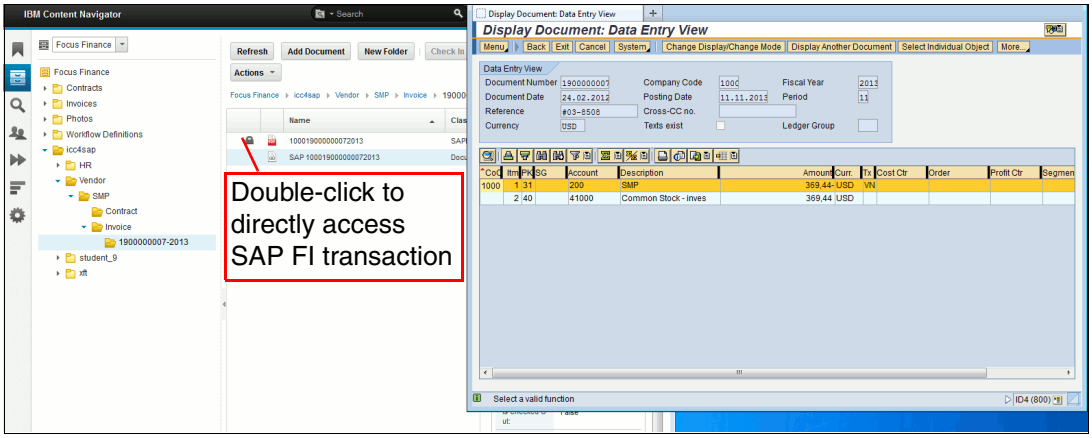

Figure 8-14 Access to the original SAP FI record from within the ECM system

Figure 8-15 shows the same information in Microsoft Office.

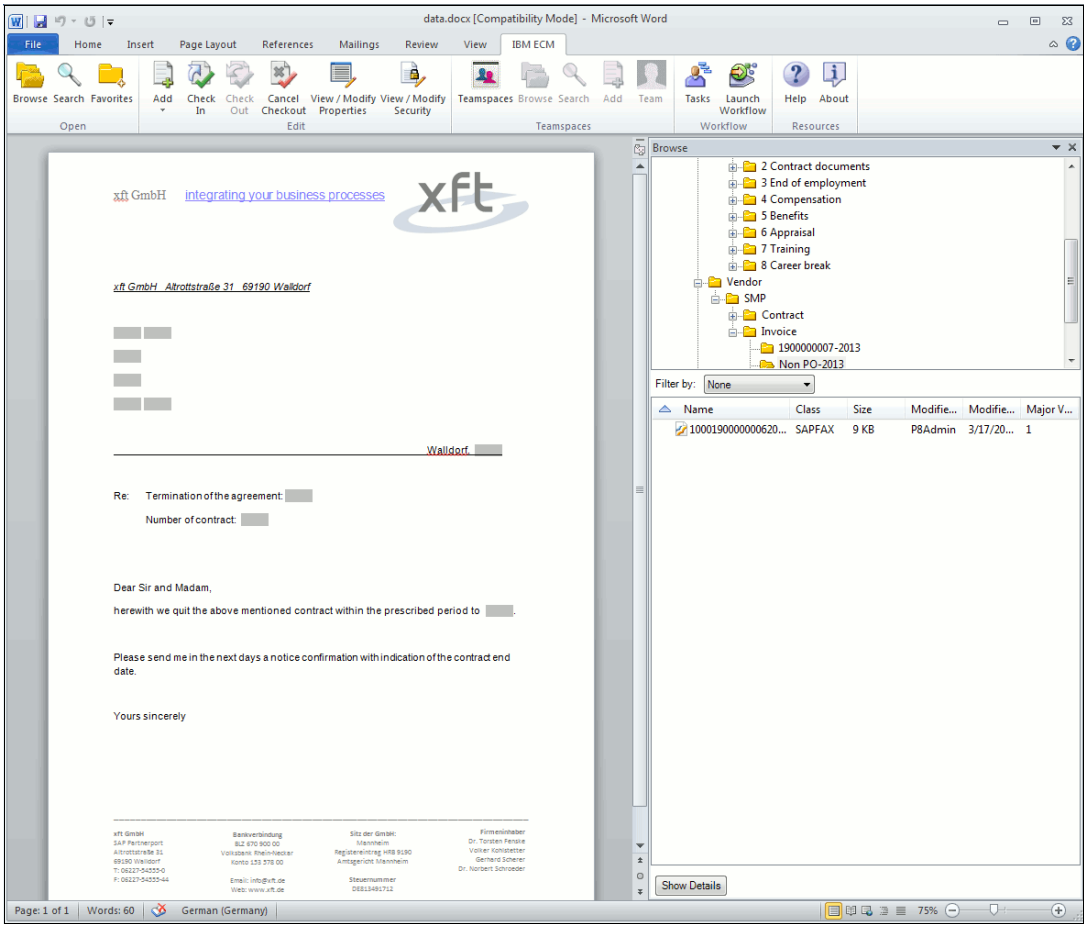

Figure 8-15 Same representation of the SAP business context in Microsoft Office

8.3.5 ECM SAP solution architecture

This section describes, as an example, the high-level architecture of the invoice processing frequently used when working with IBM and SAP Business Partners who specialize in providing SAP integrated solutions.

Figure 8-16 shows the architecture and the process flow through the components. The general use case of the integrations enabled by this architecture is to automate the capturing, processing, linking, and archiving of incoming documents to SAP processes. The example described in this section and depicted in Figure 8-16 is based on an invoice processing.

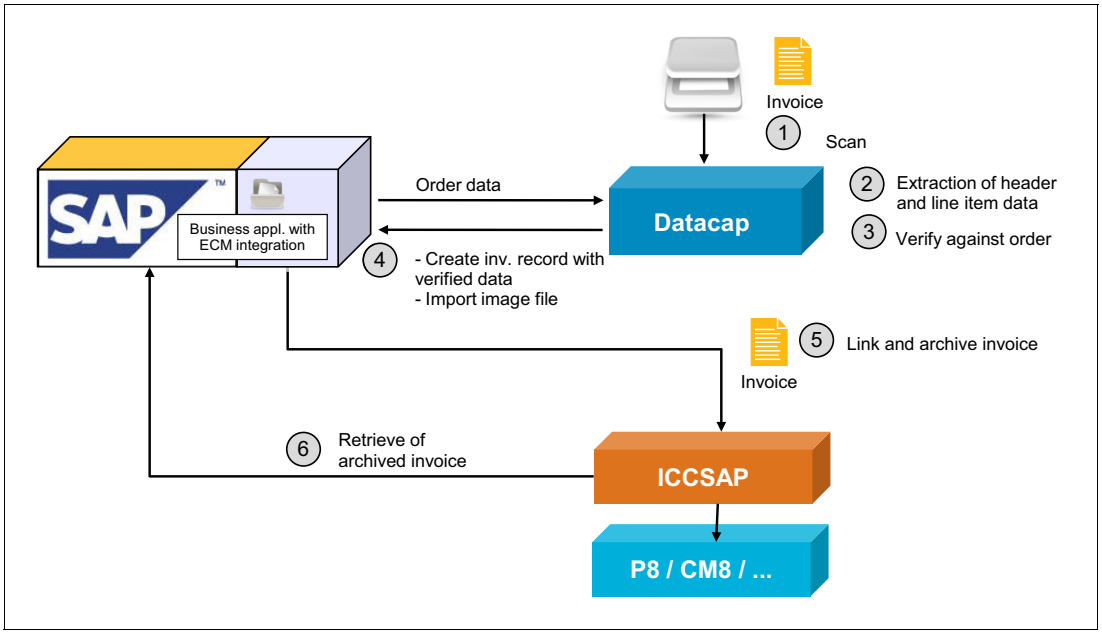

Figure 8-16 Basic architecture and example process flow

The architecture components shown in Figure 8-16 are as follows:

- IBM Datacap
- SAP ERP system (SAP ERP Central Component (ECC) 6 or later)
- ECM-centric business application, deployed on the SAP system
- IBM Content Collector for SAP Applications
- IBM repository:
  - IBM FileNet Content Manager
  - IBM Content Manager Enterprise Edition
  - IBM Content Manager OnDemand
  - IBM Tivoli Storage Manager

Figure 8-16 shows the following activities:

- Processing in IBM Datacap (steps 1 - 4)
- Processing inside the SAP business application (steps 5 - 6)

Processing in IBM Datacap

1. Scan the invoice.
2. Capture header and invoice line item data.

3. Access corresponding order data from the SAP business application and verify that the invoice line items match the order data.
4. Send the scanned invoice image and the verified invoice data to the business application in the SAP system. Create a work item in the SAP system for further processing.

Processing inside the SAP business application

5. Start and run the approval process. Archive the invoice image using SAP ArchiveLink and IBM Content Collector for SAP Applications, and link it to the work item. Post the invoice, create the invoice record, and finalize the processing.
6. The invoice document can now be retrieved as an attachment to the invoice record.

Figure 8-17 outlines the protocols and technical interfaces between the components.

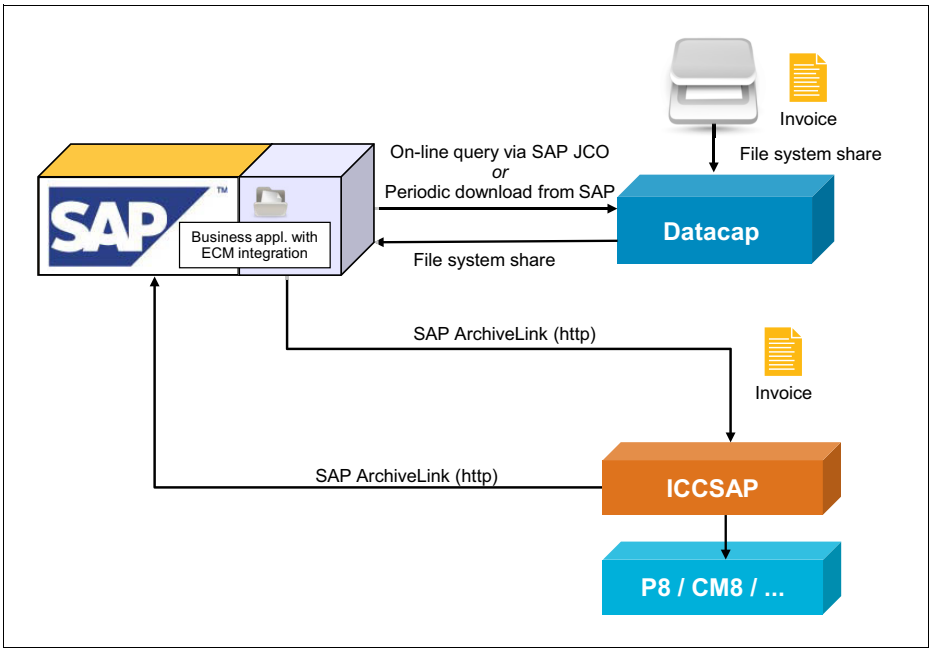

Figure 8-17 Protocols and interfaces

8.4 Data governance: Managing growth and compliance

The previous sections in this chapter extensively described the document-centric approach to ECM in an SAP context. This section focuses on the use of IBM Content Collector for SAP Applications in combination with IBM ECM repositories for SAP data archiving, following the standard SAP ArchiveLink and the SAP ILM data archiving models.

8.4.1 Business drivers

For information about how structured data (the content of the database underlying the SAP system) contributes to a large extent to the data growth that any SAP system has to maintain and control, see 8.1, “Enterprise content management business goals” on page 190. At the same time, this type of data also falls under the regulatory set of rules regarding its retention and disposal.

According to their usage patterns and their lifecycle characteristics, the following variations of structure data must be distinguished:

- Data that occurs in high volume, typically from short-lived transactions, with corresponding short-term business relevance.
- Data that falls under regulatory control because of compliance legislation, such as the Health Insurance Portability and Accountability Act (HIPAA), Sarbanes-Oxley, or other laws, typically with long-term retention requirements.
- Data that falls under the previous category of long-term retention requirements, but that resides in SAP systems that have reached, as a whole, the end of their system lifecycle.

The following sections provide more detail on each case.

High volume, short-lived transactional data

A large portion of the data stored in an SAP system originates from transactions of diverse time duration. Consider how long this data must be available instantly, as opposed to only occasionally, only in summary form, or whether it is no longer needed at all. These considerations all determine the archiving strategy.

Purchase data provides an example with a typical lifecycle, starting with the initiation of the purchase, followed by the process of an order, the execution of the order, delivery of the goods or services ordered, the billing process, and finally the payment.

After the final payment, individual order data typically only has to be available for auditing or summary reporting purposes. The usefulness of the data and the need for immediate access to it diminishes rapidly. As an element of a *closed business process*, the data is no longer of immediate importance. This type of data is ubiquitous, but in particular, and with increasing frequency, this is the case for low-cost purchases of goods and services through the Internet.

For example, the purchases of music, ebooks, and video downloads, which happen instantaneously, with delays only in the order of magnitude of minutes between the order, the delivery, the billing, and the closing of the payment process. After this point, the individual transaction is considered completely closed, it is no longer of value in the live system, and no further actions are necessary on it other than for purposes that are mostly statistical in nature.

Another part of the nature of such transactions is that they are extremely high in frequency. Therefore, they contribute massively to the data growth in the corresponding database tables that represent the transactions.

At the same time, the performance of the database in such a business scenario is of utmost importance. Therefore, a mandatory step is that preventive maintenance of the database be performed on a regular basis to preserve excellent performance for access to the database tables, and to keep transaction execution time to a minimum.

These high-frequency transactions contribute significantly to database growth, with all the associated consequences of additional cost for storage, administration, and other volume-related cost factors. Their usefulness in the database is clearly time-bound.

An operational method to address this issue is required. Obsoleted transaction data must be moved to sections of the environment that do not need to deliver high performance, and can, therefore, be provided at a lower cost. This approach does not jeopardize the still-required, albeit limited, access to the data.

Business data with regulation for long-term retention

On the other side of the lifecycle spectrum are the situations where data has long rather than short lifecycles in an organization. Industries, such as pharmaceutical companies or healthcare providers, have to comply with regulations that require their data and documents to be accessible in their original form for 20 years or longer.

Even though the active part of the lifecycle of such business data can be rather short, legal regulations increasingly require that such data remains accessible for auditing, litigation, and other documentation purposes.

Figure 8-18 shows an example of the lifecycle for business data records (and the associated attachments), with the relative importance of the data displayed as a function over the life term.

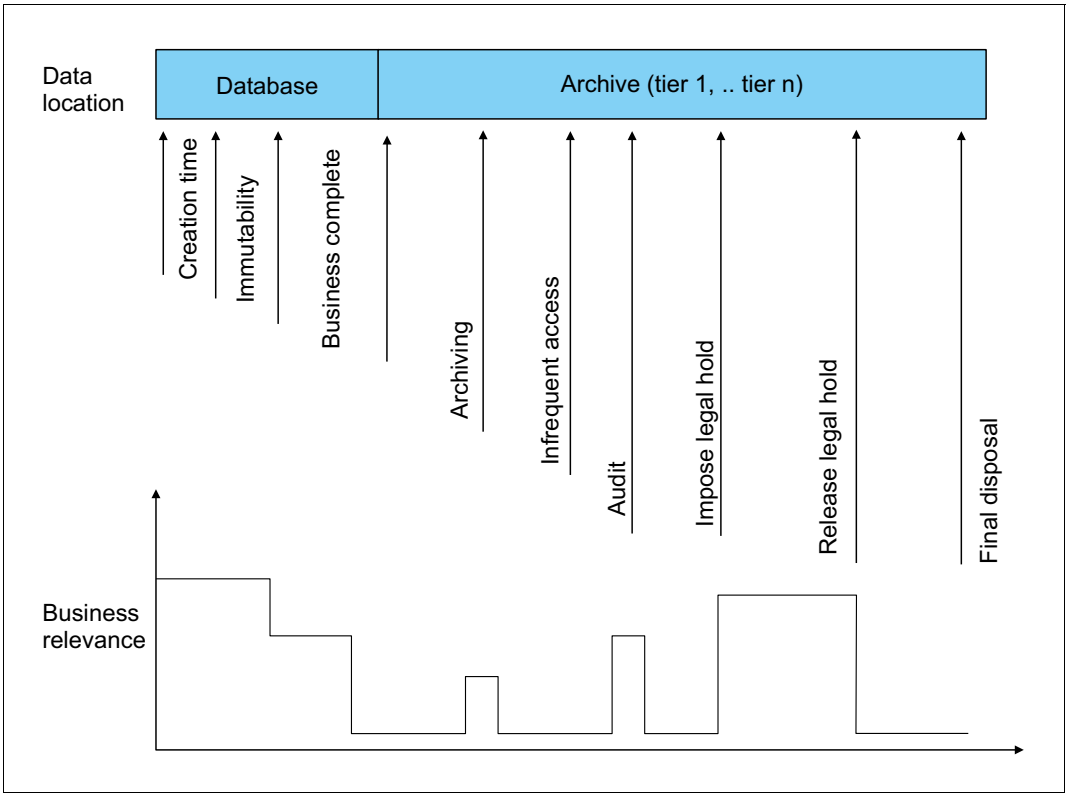

Figure 8-18 Data location and business relevance of data during the data lifecycle

Figure 8-18 shows the following data:

- Starting with creation time, and during the period where changes to the business data are still happening, importance is high, and remains so during the phase where immutability is reached, and finally the transaction can be marked as *business complete*.
- After the business complete point, the importance of the data drops significantly, until the decision is made to archive and offload the data by moving it out of the live database.
- The individual business data is still accessible, and might encounter short periods of increased relevance.
- During business audits, larger collections of archived business data are retrieved. During these periods, the data is of higher importance to support passing the audit objectives.

- The highest revival of importance that archived business data experiences is during the imposition of legal holds, typically as part of a litigation. The legal hold prevents unconditionally the disposal of data for the duration of the hold.
- After the hold has been released, the rules of retention become reactivated, and the business data eventually reaches the end of life, at which point it can (and in many cases must) be finally disposed of.

To fulfill these requirements without clogging the live SAP system, while at the same time maintaining accessibility of the data, SAP protocols for archiving and retention can be enhanced significantly through the addition of IBM middleware and secure storage solutions that ensure the immutability of the archived data.

SAP system consolidation and decommissioning

Another important core scenario for the archiving of inactive data can be found in the context of *system decommissioning*.

Typically, SAP systems are not monolithic, long-term systems in an organization-wide SAP landscape. Instead, frequently the case is that production systems are regrouped, reused, consolidated, or reorganized on a regular basis. With each of these operations, the question of data preservation as described in “Business data with regulation for long-term retention” on page 223, needs to be addressed. Not always necessary, as described previously, is to preserve all data from existing live SAP systems in its live state.

Regulations for retention and controlled destruction of business data apply to decommissioned systems just as they do for data from live systems, and therefore the triage of data described in “Leading practices for establishing data archiving” on page 227 applies equally to the decommissioning case.

All stakeholders in the business units of an organization must be closely involved in the process of deciding which business data must be preserved, over what periods of time, at what level of granularity, and when the data can or has to be securely disposed of.

Different models can be applied, depending on the amount of data that must be preserved from the system to be decommissioned, the business relevance of the data over time, and the need for immediate or delayed accessibility. The following list describes some of the options:

- Complete preservation of a system in virtualized form (made possible by current virtualization techniques).
- Transfer of the data into a consolidated new system.
- Export of data subsets into formats that can be accessed and interpreted without the need to restore into a live SAP system. The Data Retention Tool (DART) format is an example that auditors often accept as evidence, because viewers for the format are available. The viewers provide all of the audit-relevant information, without the need to keep the data in a live SAP system.

8.4.2 SAP infrastructure for data archiving

SAP provides an infrastructure with data archiving capabilities that supports the business needs outlined in 8.4.1, “Business drivers” on page 221. It fulfills the business needs by providing archiving capabilities without system interruption, and it ensures continued read access to the archived data in a transparent manner. Additionally, compression technology built into the archiving process provides additional business value by significantly reducing the cost that is directly caused by the volume of storage.

SAP data archiving is offered in the following base varieties:

- Standard data archiving using the SAP ArchiveLink protocol that is also used for document archiving. This archiving method treats the data archive as a constantly growing data store, with no explicit means to control or enforce disposal. Read access to the archived data is also achieved in a transparent manner using SAP ArchiveLink.
- Archiving in the realm of SAP ILM, with the added benefit of an explicit method to control retention periods for defined collections of data. It also provides the ability to set legal holds that override standard retention and disposal schedules for those cases where data has been put under the control of a legal dispute, and must be preserved until the dispute is closed and the hold can be released.

  ILM is based on a WebDAV model, relying on established standards for XML-based archiving, rather than on the proprietary ADK archiving format. From a business perspective, the decision to use ILM-based data archiving has to take into consideration that SAP licenses ILM separately, at additional cost, and that it requires the presence of SAP Business Warehouse.

  SAP promotes this type of archiving also in the context of system decommissioning, where data from existing systems that are no longer actively used still has to be read-accessible for compliance reasons. It should be noted that external vendors offer comparable functionality with respect to retention and legal holds capabilities, without the need to implement a full SAP ILM solution.

The following sections provide more information to help you choose between the two models based on operational and cost considerations and how they fit into the overall concept of enterprise content management using IBM ECM repositories.

8.4.3 Data archiving and the choice of IBM ECM content repositories

Structured data (pure database content) has limited value outside of the SAP system context. Its business value and its contribution to business decisions is often significant, but it is linked strictly to direct SAP operations that are executed and evaluated from within the SAP system. Operations such as content search or other types of non-SAP analytics typically do not apply to this data, unless it is extracted and transformed through other methods outside the scope of this book.

This state of affairs directly influences the choice of repository that is best suited to act as the ECM back-end for this type of archiving. The repositories can be selected on the basis of their *performance characteristics* and their *cost effectiveness*. The selection does not have to be based on advanced content use capabilities, such as content search, metadata analysis, or content aggregation.

Data archiving based on SAP ArchiveLink can be implemented with all four IBM ECM content repositories. The decision about which one to choose can be made based on the following conditions:

- Availability of existing repositories in the organization
- Archiving needs for unstructured data
- Cost considerations
- Availability of administration skills for a particular repository type

All four repository types support a tiered approach to data storage based on access performance needs, and ensure a high level of data protection.

In the most basic of possible setups, IBM Tivoli Storage Manager is often the repository of choice because of its high performance and low cost characteristics, combined with its ability to interact transparently with secure storage. For the data archiving model that is based on SAP ILM, only one IBM content repository is supported when archiving through IBM Content Collector for SAP Applications: IBM Tivoli Storage Manager.

8.4.4 SAP ArchiveLink-based data archiving

Data archiving based on SAP ArchiveLink makes use of the existing infrastructure for document archiving, and employs a customized proprietary file format for compressed storage of bundles of archived data. Data that is archived from the live SAP system is combined, bundled, and compressed into a proprietary file format: ADK or Reconciliation Object (REO). The two terms are uses in a synonymous manner. The customization of the data archiving operation must take into account the following conditions:

- The availability of SAP ArchiveLink support in the business module for which the data needs to be archived. Most SAP business modules support ArchiveLink or the Content Server interface in one form or another, and are therefore suited for this operation.
- The combination of database content into archivable objects using detailed table analysis. An important consideration is to take into account the interdependencies of SAP business objects, and how they can be combined into archivable objects that constitute a closed business transaction that can be removed from the live system.
- The applicability of retention regulations for the identified data. Even though SAP ArchiveLink has no explicit retention handling capabilities, the rules governing retention and disposal can also be applied to the stored archive files. Organizing and aggregating the data into individual files that have common properties regarding their retention policies simplifies this process.
- The scheduling of archive operations, either in manual or in automatic mode. As part of this decision point, disposal rules on an archive file basis can also be taken into account.

Depending on the business module, and its particular implementation of the archiving of its business objects, several modes of access to archived data are available.

Depending on the SAP application, archived data can either be viewed transparently from within the business transaction, or it can be interacted with separately, for example, in a document relationship browser. Note that typically this access to archived data does not imply that the data is restored from the archive back into the live system (even though this option also exists, but is only used in exceptional cases). The archive data is retrieved into the application interface for viewing purposes only.

Archiving objects: Logical units of business process elements

Data archiving enables you to move data (in compressed and standardized format) to either the local file system directly connected to the SAP system, or, through SAP ArchiveLink, to all connected content repositories.

Storing the archived data outside of the active SAP system has several advantages:

- Keep the active system lean.
- Provide advanced protection against data tampering, if, for example, certified storage systems are used.
- Improved compliance capabilities.
- Choosing an archive location physically separated from the location of the active system adds physical security and a higher level of protection for the archived data.

Associated with the archiving objects are archiving rules that determine when, and at which intervals, archiving objects will be moved from the live system to the archive. These rules can be based on temporal constraints, such as the posting date for a booking, and can be augmented with additional metadata, such as customer names or product group.

Leading practices for establishing data archiving

The need for data archiving is often ignored in the initial setup of SAP systems, and only becomes apparent when the database encounters the first wave of performance effects caused by database size. It should be a standard practice to establish a data archiving process early, rather than introducing the archiving process as a last resort before a total system performance breakdown.

This is an important point, because the setup of a well-structured data archiving system requires careful planning across multiple tiers of the organization. There must be an active interaction between the organization responsible for the operation of the SAP system (and therefore for its performance) and the SAP system business users and user groups. Also, stakeholders who are not directly related to the SAP system, but are responsible for setting governance rules and other compliance standards, must be included in the design process.

During this interaction between the interested parties, important system facts must be established:

- Expected data growth
- Detailed structure of authorizations
- Legal rules for data protection and retention times, and standard procedures to establish the scope of legal holds
- Business process interdependencies

Figure 8-19 on page 228 displays a simplified decision tree that should be applied to the data in the live system, where a triage-like sequence of decisions needs to be made based on data properties such as continued usage, the ability to summarize, and ability to remove without disruption of the live system.

A simplified decision tree that should be applied to the data in the live system (Figure 8-19).

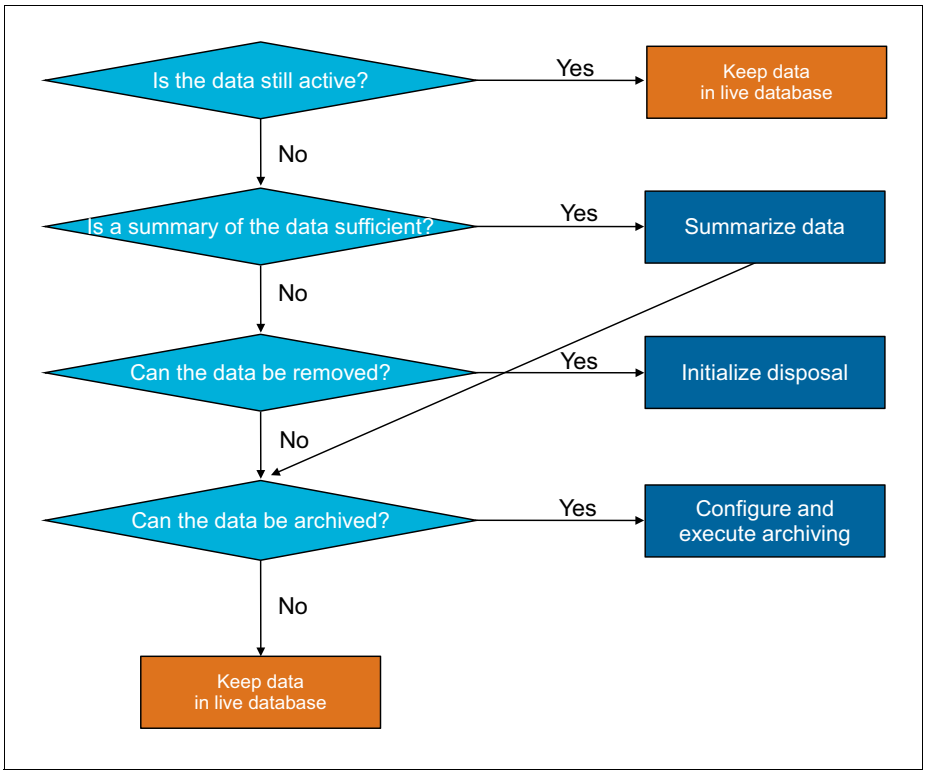

Figure 8-19 Simplified decision tree for triaging SAP data before archiving

8.4.5 Data archiving using SAP ILM

Compared to SAP ArchiveLink-based data archiving, the adoption of SAP ILM has significantly less traction in the market. This fact is mostly attributed to the complexity of the setup, and to the extra cost for the required infrastructure of such a solution.

The main use cases for data archiving using SAP ILM are as follows:

- Decommissioning of existing systems, where data from these systems still needs to be available for compliance reasons
- Business scenarios with strict retention rules or high potential of legal holds on structured data, where the retention rules can be tied to organizational structures
- Highly structured rules for identifiable subunits of the organization that map naturally into the hierarchical organization of the WebDAV-based data archive

Preparation

The implementation of ILM-based data archiving requires detailed preparation steps on the SAP side regarding the modeling of the affected data. In broader terms, this pertains to establishing a data model and retention plan alongside the business structure of the organization, reflecting the structure of rules that govern the following business aspects:

- The data that has to be retained
- The retention periods in all their variations
- Prescribed destruction times, for example, for certain types of HR data
- Definition of structural models that describe the extent of potential legal holds on structured data, based on organizational and business process modelling

The hierarchical model derived from all of these considerations defines collections of data with common properties (a property index). These properties can be used to identify groups of resources that are then assigned common retention properties.

The model provides a hierarchical, unambiguous representation of archived data, with standardized access methods using unique Uniform Resource Identifiers (URIs).

SAP ILM-enabled archive back-end

As with SAP ArchiveLink, in the ILM scenario the SAP system is also independent of the service provider for the archiving protocol. IBM provides an ILM-enabled repository that is accessed through IBM Content Collector for SAP Applications.

The intrinsic hierarchical organization of the archived data is mapped transparently into a hierarchical WebDAV archive implemented with IBM Tivoli Storage Manager. As an intermediate, IBM Content Collector for SAP Applications is used to translate the ILM WebDAV model into the IBM Tivoli Storage Manager storage model.

After this modelling has been established on the SAP side, the model has to be made known to the ILM-enabled archive (in this case the IBM Tivoli Storage Manager archive through the IBM Content Collector for SAP Applications server). It can then be used to carry out all necessary operations of archiving data, setting retention schedules at the required level of detail, applying legal holds for active litigations, and eventually disposing of the data in an auditable manner.

8.4.6 Comparison of SAP ArchiveLink and ILM-based data archiving

Which archiving model to choose in a situation depends on the business needs:

- ADK-based archiving. The main benefit is its support of database maintenance in a simple and transparent manner. It is the method of choice where none of the other requirements exist.
- ILM-based archiving. This method is the choice if complex retention and legal hold rules must be implemented that are not available, at least not internal to the protocol, for ADK-based archiving. The standardized XML-based storage format enables for more transparent access to the archived data, if such access is required from outside of the SAP system. This capability is further supported through the use of the property indexes that make the archived data transparently searchable using metadata queries.

8.4.7 Adding the value of IBM middleware and storage solutions for SAP data archiving purposes

In summary, both data archiving protocols of SAP support the implementation of external archive providers using standardized interfaces. IBM ECM products provide the necessary flexibility regarding the choice of repository. IBM storage solutions ensure the required security conditions to achieve full compliance with the imposed regulations.

All IBM ECM repositories provide storage hierarchy solutions that will store the data in the most cost-effective way based on usage patterns and accessibility requirements. Full integration into the IBM ECM product portfolio enables the archiving solution to extend beyond the base requirements of the SAP archiving standards into complete end-to-end solutions.

8.5 References

These websites are also relevant as further information sources:

- IBM enterprise content management portfolio
  http://www.ibm.com/software/products/en/category/enterprise-content-management
- IBM Information Lifecycle Governance solutions
  http://www.ibm.com/software/products/en/category/information-lifecycle-governance
- IBM Value-Based Archiving solutions
  http://www.ibm.com/software/products/en/value-based-archiving
- IBM Datacap
  http://www.ibm.com/software/info/datacap/
- IBM Content Collector for SAP Applications
  http://www.ibm.com/software/products/en/content-collector-sap
- IBM Content Collector for SAP Applications Knowledge Center
  http://www.ibm.com/support/knowledgecenter/SSRW2R_4.0.0
- IBM Content Collector for SAP Applications: Sizing, Configuration, and High Availability
  http://www.ibm.com/support/docview.wss?uid=swg27036773

9

IBM Business Analytics infrastructure for SAP

This chapter provides architecture guidelines to efficiently integrate IBM business analytics (BA) software with SAP solutions to gain an end-to-end analytical perspective on heterogeneous enterprise data, and build a sustainable IBM-SAP business analytics solution landscape.

With the rise of big data, data assets have become larger, appear in many more variations, and arrive at higher speeds in an enterprise. Business analytics for making informed decisions yielding better business outcomes must be able to process all of the data in all its forms (from structured to unstructured), either at rest or in motion. Not surprisingly, many specialized solutions are available for different analytics needs.

Comprehensive coverage of all analytic solutions is beyond the scope of this book. To narrow the scope, the following topics are *not* included in this chapter, but they are extensively described in other IBM publications, which are listed in 9.5, “References” on page 248:

- Hadoop-based analytics. Data at rest, including the full spectrum from well-structured (relational) to least-structured (video, text, audio, and so on) data.
- Streaming analytics. Data in motion, including the full spectrum from well-structured (relational) to least-structured (video, text, audio, and so on) data.
- Decision support systems. Data in motion for operational decision making in real time, for next best offer, next best action, and so on.
- IBM Content Analytics with Enterprise Search, such as click-stream analytics.
- Cognitive computing with IBM Watson™.

This chapter includes the following topics:

- 9.1, “IBM Business Analytics infrastructure for SAP value proposition” on page 232
- 9.2, “IBM Business Analytics integration architectures” on page 237
- 9.3, “Detailed review of IBM Business Analytics integration architectures for SAP” on page 239
- 9.4, “Conclusion” on page 247
- 9.5, “References” on page 248

© Copyright IBM Corp. 2014, 2015. All rights reserved.

9.1 IBM Business Analytics infrastructure for SAP value proposition

Enterprise resource planning (ERP) systems, such as the SAP system, provide a highly robust transactional environment that is typically optimized for transactional speed and configurability of core business functionality. ERP systems, such as SAP, excel at managing transactions and day-to-day business. Clients often consider the cost of deploying ERP systems to be a part of the cost of doing business.

Alternatively, business analytics solutions excel at enabling better business decisions to manage and drive business performance by providing complete visibility and fast insights into the business. IBM Business Analytics software can increase the value of SAP investments by driving new business insights from SAP data, especially when combined with non-SAP data in a heterogeneous enterprise.

Some organizations, in which users need to perform swift analysis of data from SAP solutions, require a faster-performing business intelligence (BI) solution. SAP has introduced their SAP in-memory appliance, SAP High-Performance Analytic Appliance (HANA) as a solution to help organizations analyze large volumes of detailed operational and transactional information in real time.

IBM Business Analytics infrastructure for SAP provides near real-time in-memory analytic capabilities for SAP and non-SAP applications that do not require investment in HANA. However, if HANA is already a part of the enterprise landscape, IBM Business Analytics infrastructure for SAP provides seamless integration with HANA, as described later in this chapter.

IBM Global Technology Outlook 2013[1] highlights how recent advances in technology are changing the relationship between an organization and the people it interacts with, whether they are customers, business partners, or employees. It describes the emergence of the *contextual enterprise*.

A contextual enterprise is an organization that dynamically adapts to the changing needs of its individual customers by using information from a wide range of sources. It uses information from multiple sources to provide meaningful context to communications with customers, business partners, and analysts. Contextual enterprise requires an architectural perspective to eliminate a siloed approach to enterprise analytics, and to augment existing systems with new information to support requirements that are needed by the contextual enterprise.

This information architectural approach has several goals:

- Combine the best information from across the organization, for use both by operational and analytical systems.
- Augment this consolidated information with relevant facts and event triggers, by analyzing previously untapped sources of information.
- Help to more easily locate information and deliver it to where the business most needs it.

In a heterogeneous IT landscape, SAP systems are an important (but not the only) source of information for contextual enterprise analytics. SAP data must be combined with business data from other enterprise systems to enable contextual enterprise.

[1] IBM Global Business Outlook 2013: http://www.zurich.ibm.com/pdf/isl/infoportal/Global_Technology_Outlook_2013.pdf

The next generation business analytic solutions from IBM help organizations of all sizes make sense of information in the context of their business. Organizations can uncover insights more quickly and more easily from all types of data, even big data, and on multiple platforms and devices.

In addition, with self-service and built-in expertise and intelligence, organizations have the capabilities and confidence to make smarter decisions that better address their business imperatives. IBM offers flexible deployment options for business analytics solutions, including software as a service (SaaS) options, to mitigate concerns regarding costs and the complexity of deployment.

IBM Cognos software helps organizations realize a greater return on their investments in SAP applications with faster access to the data that the business needs to make smarter decisions. When IBM Cognos software is integrated with SAP applications, the value of SAP data is enhanced, and users gain the perspective and context needed to derive insight from SAP data.

In addition, using IBM Cognos software and SAP applications together can help minimize the number of tools and duplicate content that organizations must maintain, streamline training requirements, and significantly reduce IT backlogs. Cognos is a data source-independent, best-in-class business analytics platform well suited for a heterogeneous enterprise, while providing an extensive set of SAP-certified integration services.

IBM Cognos Business Intelligence brings together reporting, analysis, scorecarding, and dashboards. It expands these BI capabilities with planning, scenario modeling, real-time monitoring, and predictive analytics.

Cognos Business Intelligence enables organizations to access information within the organization and beyond, to connect to key stakeholders, and to share insight, align, and make decisions. In so doing, IBM Cognos Business Intelligence unleashes the collective intelligence in the organization, so you can “see around corners”, predict outcomes, make informed decisions, and act smarter and faster than the competition.

Cognos Business Intelligence enables organizations to accomplish the following tasks:

- Equip users with the tools that they need to explore information freely, analyze key facts, collaborate to gain alignment with key stakeholders, and make decisions with confidence for better business outcomes.
- Provide quick access to facts with reports, analysis, dashboards, scorecards, planning, budgets, real-time information, statistics, and the flexibility to manage information for more informed decisions.
- Integrate the results of *what-if* analysis modeling and predictive analytics into a unified workspace to view possible future outcomes alongside current and historical data.
- Support wherever users need to work with BI capabilities for the office and desktop, on mobile devices, online, and offline.
- Meet different analytics needs throughout the business with solutions that are integrated and right-sized for individuals, workgroups, or midsize businesses and large organizations or enterprises.
- Implement a highly scalable and extensible solution that can adapt to the changing needs of IT and the business, with flexible deployment options that include the cloud, mainframes, and data warehousing appliances.
- Start to address the most pressing needs of the organization with the confidence that the solution can grow over time to meet future requirements with the integrated IBM Cognos family of products.

IBM Cognos Mobile extends interactive Cognos Business Intelligence to a broad range of mobile devices, including the Apple iPhone and iPad, Android, and tablets. With a rich client, users can view and fully interact with Cognos reports, dashboards, metrics, analysis, and other information in a security-rich environment. Users receive timely, informative and interactive BI to support their decision making, regardless of location.

Beyond BI capabilities, business managers today need greater analytical ability to manage their business performance. They require a view to analyze data from SAP applications, but they also need to forecast variations in revenues and expenses, follow customer trends, uncover drivers of hidden costs, and respond to economic or competitive changes whenever the need arises.

IBM Cognos TM1 is a market-leading enterprise planning software that enables organizations to collaborate on plans, budgets, and forecasts. Cognos TM1 enables users to analyze data and create models, including profitability models, to reflect a constantly evolving business environment. In addition, integrated scorecards and strategy management capabilities help organizations monitor performance metrics and align resources and initiatives with corporate objectives and market events.

TM1 capabilities for ad hoc analysis, scenario modeling, and collaborative forecasting extend and complement transactional operations performed by ERP systems. TM1 software easily accesses data from SAP Business Warehouse (BW) or SAP ERP Central Component (ECC), and organizes complex business information so that organizations can evaluate current and past performance, perform *what-if* analysis, and forecast resources in real time to consider future scenarios.

TM1 features memory-based, multi-dimensional cube architecture. The online analytical processing (OLAP) engine driving TM1 yields excellent response times. In addition, with multiple memory-based cubes, data is more rapidly searched, modified, and restructured than with a single-cube, disk-based structure.

Predictive analytics helps organizations to use all available data, and predict with confidence what will happen next, so that you can make smarter decisions and improve business outcomes. IBM offers easy-to-use predictive analytics products and solutions, such as IBM SPSS, that can use data from SAP and non-SAP sources and meet the specific needs of different users and skill levels, from beginners to experienced analysts.

With predictive analytics software from IBM, organizations can achieve the following goals:

- Transform data into predictive insights to guide front-line decisions and interactions.
- Predict what customers want and will do next to increase profitability and retention.
- Maximize the productivity of their people, processes, and assets.
- Detect and prevent threats and fraud before they affect the organization.
- Measure the social media effect of their products, services, and marketing campaigns.
- Perform statistical analysis, including regression analysis, cluster analysis, and correlation analysis.

Integration between IBM SPSS Modeler and SAP is addressed in 9.3.5, “Predictive analytics with SAP” on page 245.

9.1.1 Architecture overview

The reference architecture for IBM Business Analytics infrastructure for SAP is built on the common enterprise data warehouse (EDW) that is used for both SAP and non-SAP enterprise data (see Figure 9-1).

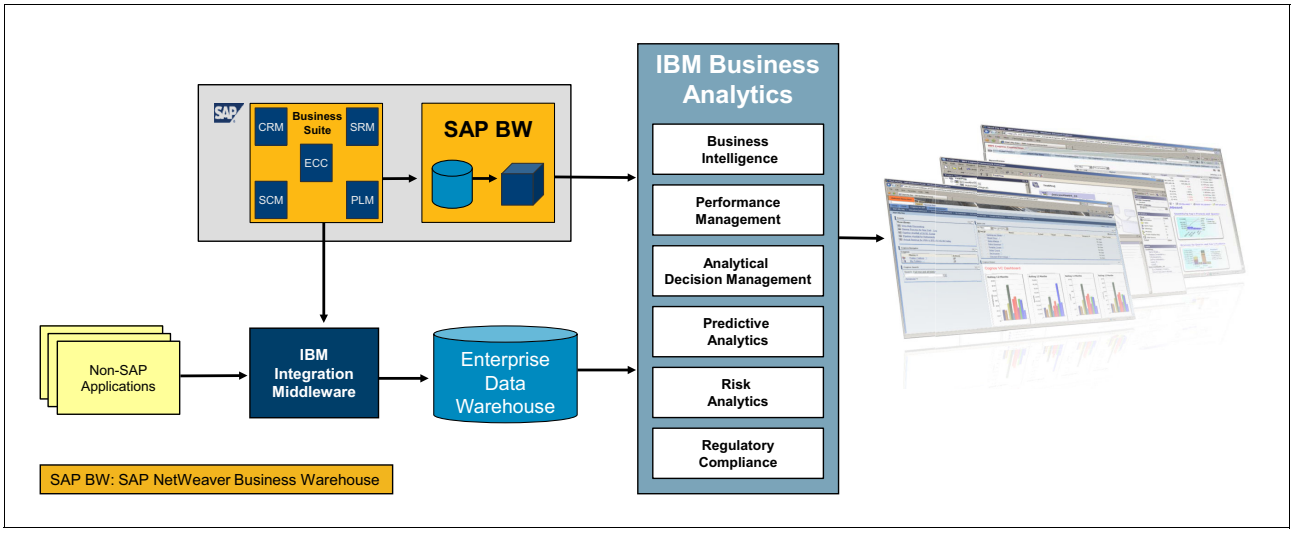

Figure 9-1 Reference architecture for IBM Business Analytics infrastructure for SAP

This reference architecture uses SAP BW as a packaged solution. SAP BW is an analytical, reporting, and data warehousing solution produced by SAP. SAP BW is a packaged solution that includes SAP delivered extractors, data models, and queries (also known as Business Content). In this reference architecture, the scope of SAP BW is limited to operational reporting on SAP Business Suite data. All deeper analytics and all analytics at an enterprise level should be based on the EDW.

A key architectural decision in this approach is to use SAP Business Suite directly as the source of data for analysis, as opposed to extracting SAP data from SAP BW. The reference architecture described in this chapter takes into account the following considerations:

- SAP BW typically does not have all of the detailed data needed for analytics required to support the contextual enterprise.
- IBM data extraction from SAP BW might require additional SAP licensing (for SAP Open Hub).
- IBM has built integration capabilities that make it easier to extract data from SAP Business Suite applications into EDWs.

IBM Business Analytics offerings and capabilities shown on Figure 9-1 provide seamless integration with SAP BW.

The integration of IBM Business Analytics solutions with SAP solutions can be fundamentally of the following kinds:

- Based on analysis of SAP data exported into a persistent data store outside of SAP
- Direct access to SAP data

In an IBM Business Analytics solution for SAP, a critical step is to ensure that the middleware perspective is aligned with the core BA architecture principles in the following ways:

- Open. Robust and standardized extract, transform, and load (ETL) capability to source, transform, and load all types of data, including flat file, relational databases, dimensional databases, unstructured, and structured data.
- Integrated. End-to-end business processes with minimum bridges and interfaces to mitigate risks of data inconsistency.
- Optimized for information quality and lifecycle management. Data cleansing policies, data retention policies, archiving, and near-line storage.
- Technology alignment. Promote pre-delivered packages aligned with innovative product roadmaps from SAP and IBM.
- Minimize total cost of ownership (TCO). Reduce complexity, for example, on interfaces, maintenance, stability, and skill set.

Figure 9-2 provides an overview of IBM Business Analytics integration capabilities for SAP.

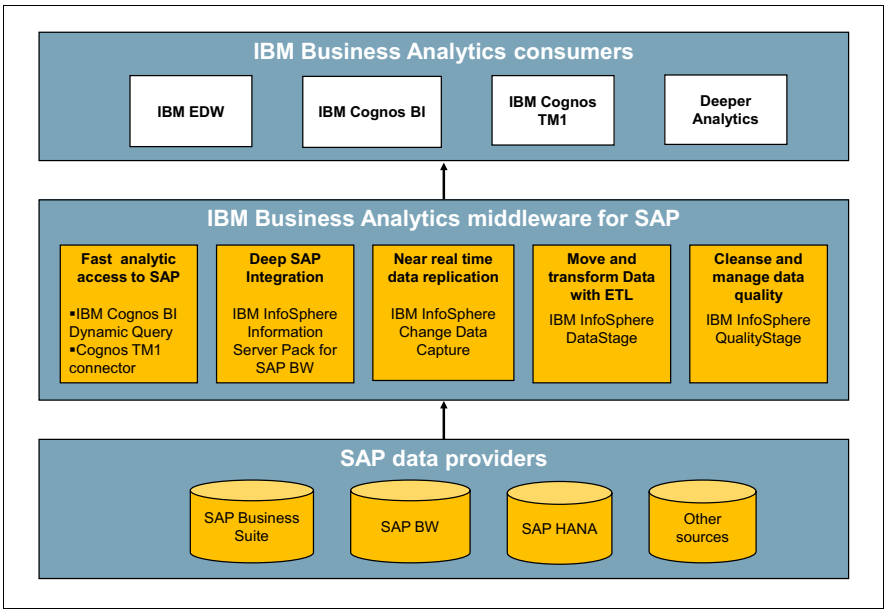

Figure 9-2 Overview of IBM Business Analytics integration capabilities for SAP

Figure 9-2 shows that IBM middleware has capabilities accomplish the following tasks:

- Connect to any SAP source.
- Deeply integrate with service-oriented architecture (SOA) data.
- Provide near real-time SAP data replication into analytics solutions.
- Provide structured data extraction and cleansing capabilities.

The IBM InfoSphere Information Server is a key component that encapsulates best-in-class integration tools to collect metadata, and to manipulate or assess data before integration with consumer BA applications. SAP integration is based on using SAP-certified integration interfaces:

- Open Database Connectivity (ODBC)
- Java Database Connectivity (JDBC)
- OLAP Business Application Programming Interface (BAPI)
- Remote Function Call (RFC) through Advanced Business Application Programming (ABAP) business functions

The following sections describe a set of integration patterns for IBM Business Analytics infrastructure for SAP, and provide selection guidelines for, and the corresponding benefits of, each option.

9.2 IBM Business Analytics integration architectures

This section describes the various integration architectures for integrating SAP source systems with IBM Business Analytics consumers.

Figure 9-3 shows several integration architectures for key SAP data providers and IBM consumers. These architectures are based on either retrieving SAP data in a streaming mode (*direct access*) or requiring SAP data to be extracted into non-SAP persistent data storage (*data export*).

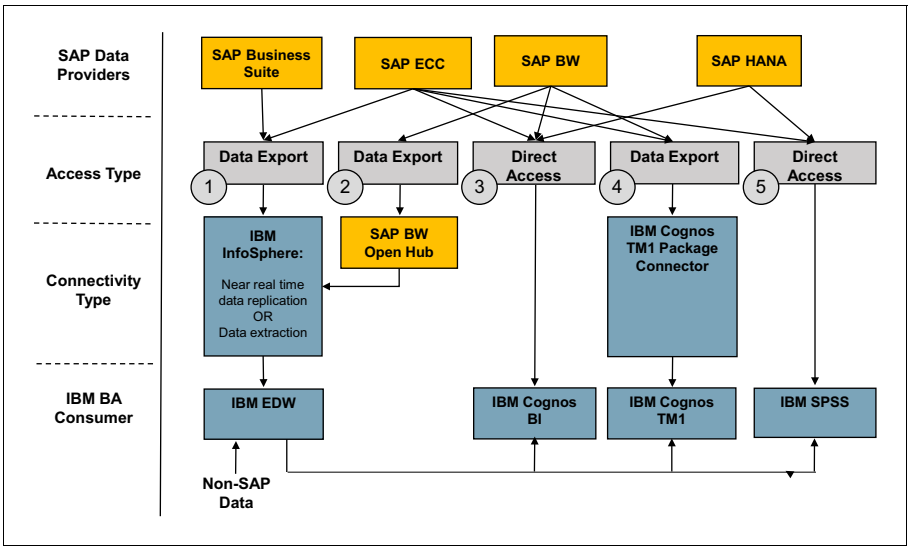

Figure 9-3 Integration architectures for IBM Business Analytics infrastructure for SAP

The following list describes the integration architectures shown in Figure 9-3:

1. This integration architecture is based on exporting data from SAP ECC and other applications from SAP Business Suite into an IBM EDW. A data warehouse is a system that enables you to separate online transaction processing (OLTP) data used by business applications to record business transactions from data needed for decision-support systems (OLAP).

   Data export from SAP into EDW is implemented by IBM middleware, in this case, IBM InfoSphere. In the EDW, SAP data is combined with non-SAP enterprise data from other data sources. Subsequently, EDW is used by Cognos Business Intelligence, TM1, and SPSS tools for analytical purposes.

2. This integration architecture supports data extraction from SAP BW into an IBM EDW using the same IBM InfoSphere middleware. This architecture requires the SAP middleware component, SAP BW Open Hub.

3. This integration architecture is based on the direct connectivity of Cognos Business Intelligence tools to various SAP source systems. In this case, no data extraction takes place.

4. This integration architecture enables you to extract SAP data into TM1 in-memory cube database, and subsequently conduct various kinds of planning, scenario modeling, and many other types of analytics.
5. This integration architecture is based on direct connectivity of IBM predictive analytics tools to SAP source systems.

For more details about these architectures, see 9.3, “Detailed review of IBM Business Analytics integration architectures for SAP” on page 239.

9.2.1 IBM Enterprise Data Warehouse products

This section describes several IBM products which can be used to implement EDW.

IBM DB2 with BLU Acceleration

IBM DB2 with BLU Acceleration speeds analytics and reporting using dynamic in-memory columnar technologies. In-memory columnar technologies provide an efficient way to scan and find relevant data. Coupled with other innovations, such as parallel vector processing and actionable compression, it makes the analytics queries far faster and less complex. BLU Acceleration is a collection of technologies developed by IBM and integrated directly into the DB2 engine.

BLU Acceleration is a new storage engine, along with integrated run time, that supports the storage and analysis of column-organized tables. The BLU Acceleration processing is parallel to the regular, row-based table processing found in the DB2 engine. BLU Acceleration is not a bolt-on technology or a separate analytics engine that sits outside of DB2. Row-based table storage and column-based table storage both coexist in the same database and Structured Query Language (SQL) can access them both at the same time.

DB2 with BLU Acceleration does not require all data to be in random access memory (RAM). Because of a combination of software technology, such as active compression, columnar compression, scan-friendly caching, and data skipping, only a fraction of the data needs to be buffered in RAM. DB2 BLU is *memory optimized*, designed to deliver massive performance improvements even when only a small fraction of data can fit into the memory resources.

IBM PureData System for Analytics

IBM PureData® System for Analytics is the next generation of the IBM Netezza® Data Warehouse Appliance product. It has the same key design tenets of simplicity, speed, scalability, and analytics power that was fundamental to Netezza appliances. It combines the large data processing efficiency based on massively parallel processing (MPP) architecture with hardware-level SQL acceleration.

With simple deployment, ready-to-use optimization, no tuning requirements, and minimal ongoing maintenance, PureData System for Analytics leads the industry with fast time-to-value and low TCO.

IBM PureData System for Operational Analytics

IBM PureData System for Operational Analytics is an expert, integrated data system designed for operational analytics workloads across the enterprise. The system offers both the simplicity of an appliance and the flexibility of a custom solution. Designed for high performance, it can handle up to 1000 concurrent operational queries.

IBM PureData System for Operational Analytics provides the following capabilities:

- Fast performance using parallel processing technology and other advanced capabilities
- Built-in expertise and analytics to help you expertly manage database workloads at lower cost
- Simpler administration for easier management and lower cost of ownership

IBM DB2 Analytics Accelerator for z/OS

IBM DB2 Analytics Accelerator for z/OS® is a high-performance appliance that integrates IBM Netezza and IBM zEnterprise® technologies. The solution delivers extremely fast results for complex and data-intensive DB2 queries on data warehousing, BI, and analytic workloads.

9.2.2 IBM InfoSphere DataStage

IBM InfoSphere DataStage is the premium tool to extract, transform, and load data from any source system into an EDW. It offers parallel processing for loading large volumes of data in the shortest time possible. It also features an intuitive visual interface to design complex business rules, known as *business flows*, and includes key capabilities for error handling and monitoring.

9.2.3 IBM InfoSphere Data Replication

For real-time business requirements, the IBM InfoSphere Information Server Family provides the IBM InfoSphere Data Replication tool that supports low-impact change data capture at the database level. In addition to real-time replication, this technology enables delta or incremental loads, which significantly reduces the data replication time windows, and mitigates data replication performance issues.

9.3 Detailed review of IBM Business Analytics integration architectures for SAP

This section provides details of the Business Analytics integration architectures shown in Figure 9-3 on page 237, including components and technical prerequisites.

9.3.1 Data export from SAP Business Suite into an IBM enterprise data warehouse

This integration architecture enables centralized enterprise decision-making processes, and drives business performance by providing complete visibility and fast insights into the business. IBM EDW uses integrated data from any SAP Business Suite system and data from non-SAP enterprise systems, and provides an information asset to support analytics and the decision-making processes.

The architecture shown in Figure 9-4 on page 240 uses DataStage. SAP business data is extracted from SAP Business Suite into an IBM EDW using SAP integration provided by IBM InfoSphere Information Server Pack for SAP Applications.

When the target EDW implementation is PureData for Analytics (formerly known as IBM Netezza Data Warehouse Appliance), the data export can be a simple DataStage job that copies data from a set of SAP tables to the target. In this case, because of extreme

performance of SQL queries in PureData for Analytics, it is possible to bypass the data transformation into a traditional format used by data warehouses (*star* schema).

For other EDW implementations, data transformation into a *star* schema can be advantageous, and is fully supported by standard ETL techniques, which can include staging tables and data cleansing, as shown in the lower part of Figure 9-4.

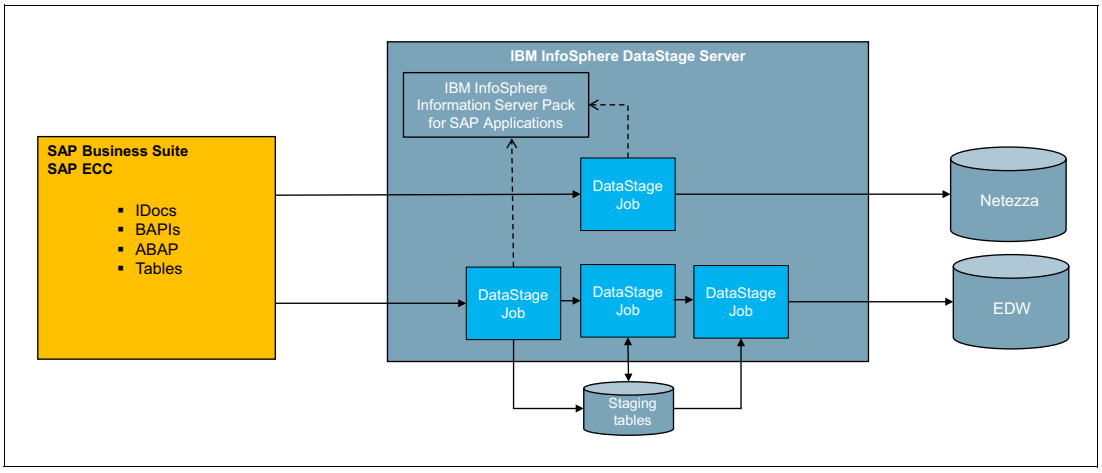

Figure 9-4 Data export from SAP Business Suite into an IBM EDW

This integration supports the following SAP interfaces:

- SAP tables (cluster or pooled tables). This integration uses IBM integration tools to discover SAP tables by accessing SAP data dictionaries, and automatically generate and deploy RFC-enabled ABAP code modules into the SAP system. These modules are subsequently used for extracting data from SAP tables.
- Intermediate Documents (IDocs) in SAP proprietary format as part of SAP Application Link Enabling (ALE) interface. ALE is an integration technology that can integrate business processes between SAP systems and non-SAP systems, and between SAP systems.
- BAPIs. BAPI is a precisely defined interface providing access to processes and data in SAP business application systems.
- ABAP business functions that can be called remotely using the RFC mechanism. This is the same mechanism that is used to call BAPIs, but this category of remotely accessible SAP programs includes any RFC-enabled function modules, not just well-known and documented BAPI functions.

This architecture offers the following benefits:

- Provides a cost-effective option to accelerate time-to-market, because most of the procedures are automated, and minimum setup and configuration is required
- Mitigates data inconsistency risks with DataStage error handling and monitoring
- Provides continuous data changes through ETL staging tables, which improves the overall data visibility for lines of business (LOBs)

9.3.2 Data export from SAP BW into an IBM EDW

In a heterogeneous enterprise landscape, multiple local, departmental, and regional data warehouses or data marts might need to coexist and be used to feed an EDW with aggregated data. One or more of such local, departmental, and regional data warehouses can be an SAP BW system. For instance, consider an example of a country-specific

SAP HR analytics system feeding a data warehouse for consolidated HR reporting at the corporate level.

This integration architecture, shown in Figure 9-5, enables you to pull data from SAP BW into an EDW. Subsequently, IBM BA tools, such as Cognos Business Intelligence or SPSS, can be used to conduct BA tasks from the EDW.

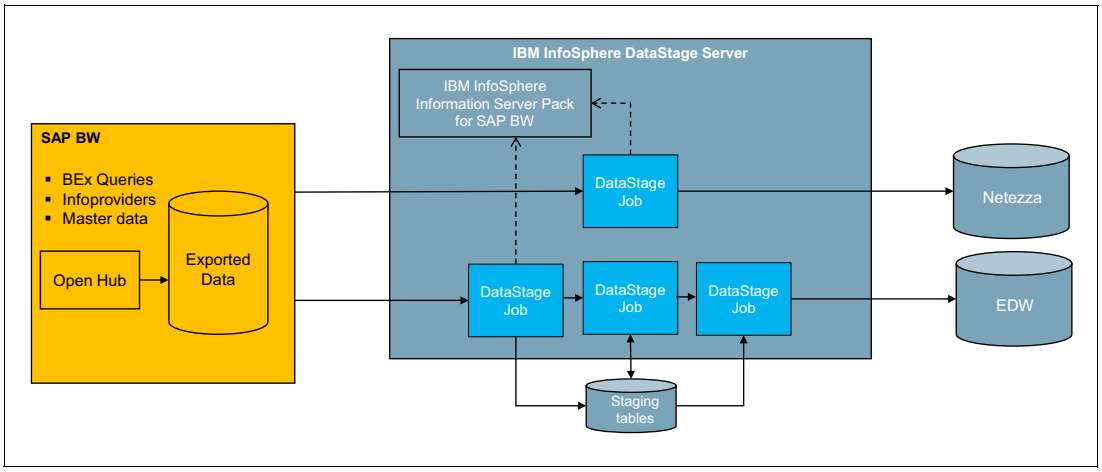

Figure 9-5 Data export from SAP BW into an IBM EDW

SAP provides only one supported mechanism to extract data from SAP BW to non-SAP repositories, which is an SAP BW Open Hub extract capability. With the Open Hub service, you can model, schedule, run, and monitor the data export of various data entities in SAP BW, such as Business Explorer (BEx) queries, InfoCubes, master data, and so on, into, for example, a set of destination database tables inside SAP BW.

The database tables in SAP BW can then be used by external consumers to get data out of SAP. Open Hub is integrated into SAP BW, but typically requires additional licensing from SAP. Subsequent data flow in this scenario might require complex data validations or lookup rules. DataStage provides developers with an intuitive development platform to build complex data flows.

A staging table, as shown in Figure 9-5, temporarily hosts the data extracted from SAP BW through an Open Hub. To implement the extraction from the DataStage server, create a process chain in SAP BW and include a process step `Open hub destination` execution. The IBM InfoSphere DataStage Open Hub Extract job then triggers the execution of the process chain, and initiates the inbound data flow through the DataStage server into the IBM target consumer, which can be an IBM data warehouse or an IBM database.

IBM InfoSphere Information Server Pack for SAP BW reduces the project time and cost of distributing and integrating data from SAP BW. It supports extracting information from SAP BW for use in other data marts, data warehouses, reporting applications, and other targets. Using SAP BW Open Hub services, the extract capability enables users to graphically browse and select Open Hub targets. InfoSphere Information Server Pack for SAP BW provides SAP-certified integration for both loading and extracting data, including Unicode.

InfoSphere Information Server Pack for SAP BW also provides direct access to, and creation of, SAP BW metadata within the DataStage user interface (UI). Users can browse, select, create, and change SAP BW metadata objects (Source Systems, InfoSources, InfoObjects, InfoCatalogs, and InfoPackages) using complete metadata integration capabilities.

Similar to the integration architecture described in 9.3.1, “Data export from SAP Business Suite into an IBM enterprise data warehouse” on page 239, when EDW implementation is in PureData for Analytics, the data export can be a simple DataStage job that copies data from a set of SAP tables to the target. The extreme performance of SQL queries in PureData for Analytics makes it possible, which enables you to bypass the data transformation into formats optimized for data warehouses.

9.3.3 Operational analytics with Cognos Business Intelligence directly accessing SAP solutions

This integration architecture uses Cognos Business Intelligence, an enterprise-standard BI reporting tool set, as a self-service platform for business users to perform operational analytics for SAP. Figure 9-6 shows this integration architecture. SAP data is consumed on demand directly from SAP ECC, SAP BW, and SAP HANA. Reports are generated based on Cognos Business Intelligence tools from the SAP source system.

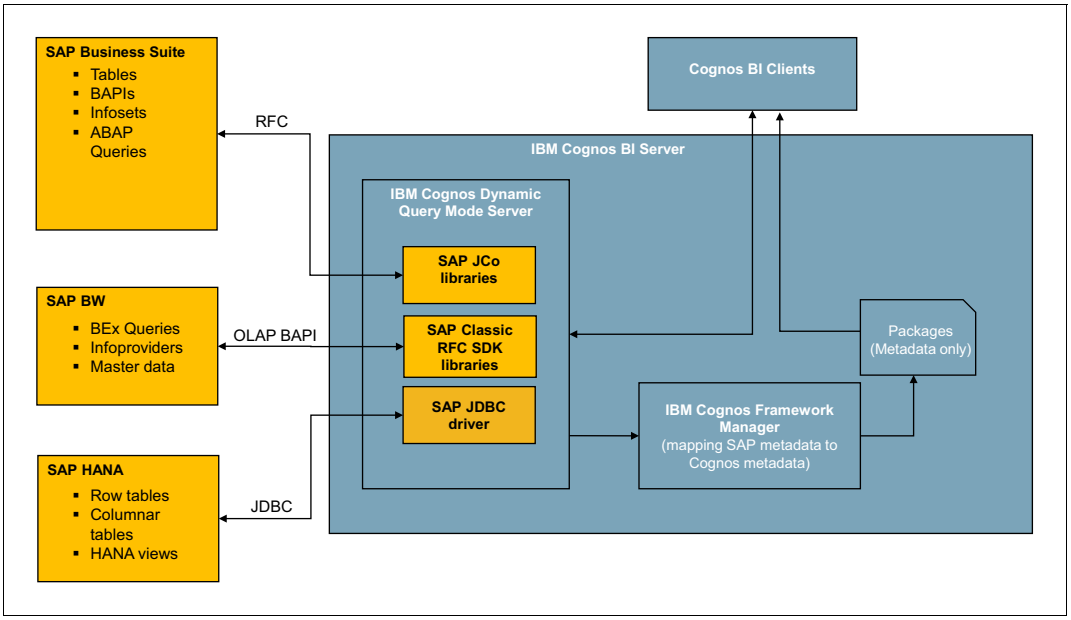

Figure 9-6 Operational analytics with IBM Cognos Business Intelligence directly accessing SAP

SAP ECC data can be accessed using SAP tables, BAPIs, SAP Infosets, and ABAP queries. An Infoset is a special view of a data source (list of fields). It is the basis of an ABAP query, which represents a selection of data from an Infoset. This approach enables you, for example, to generate operational reports and calculate key performance indicators (KPIs) based on granular data, and on a daily basis.

SAP BW data providers can be BEx queries, InfoProviders, or master data (text, attributes, or hierarchies). The term *InfoProvider* encompasses objects that physically contain data, for example, InfoCubes and DataStore objects. InfoProviders can also be objects that do not physically store data, but that display logical views of data, such as VirtualProviders, InfoSets, and MultiProviders. BEx queries are a preferred choice as data providers, because you can define data restrictions based on key analysis dimensions to streamline the data bandwidth.

SAP HANA provides standard interfaces to existing applications, operational software, and other business applications. It enables organizations to use investments in existing BI clients (including Cognos Business Intelligence) for access to the information available in SAP HANA systems.

By applying this integration scenario, you can stream SAP data from any SAP HANA database (row or columnar tables and HANA views), and consume it through any Cognos Business Intelligence Clients. The technical prerequisite is to have the SAP JDBC driver installed on the IBM Dynamic Query Mode Cognos Business Intelligence server. Note that the Change Data Capture (CDC) feature is not supported in this case (SAP HANA as the source system).

Except for the import of metadata managed through IBM Cognos Framework Manager, no SAP data storage persistency occurs at the IBM middleware level (IBM Cognos Dynamic Query Mode server or IBM Cognos Business Intelligence server). Dynamic Query Mode provides fast analysis capabilities by using in-memory technology to cache data result sets from SAP in the Cognos Business Intelligence server.

This in-memory technology provides an enhanced Java-based query mode that offers several key capabilities:

- Query optimizations to simplify and speed up queries, and reduce data volumes with improved query execution techniques
- Significant improvement of complex OLAP queries through intelligent combinations of local and remote processing, and better Multidimensional Expression Language (MDX) generation
- Security-aware caching with 64-bit processing

The connection between the SAP source system and Cognos Business Intelligence server is based on RFC calls. The SAP Java connector (SAP JCo) library is installed on the Cognos Business Intelligence server. The preferred approach in this scenario is to choose an ABAP query as the data provider, to reduce the data result set to be streamed from an end-to-end perspective. ABAP query is essentially an SAP report object generated using SAP tools, which avoids the need for ABAP coding.

The following list describes the benefits of this architecture:

- Quick setup and configuration with the Cognos Dynamic Query Mode server acting as a gateway between the SAP Business Suite source system and the Cognos Business Intelligence server
- Use of SAP BAPIs or InfoSets pre-built content for data sourcing (ready-to-use solution, including complex business logic)
- Connectivity options to have BAPIs, InfoSets, or ABAP queries as data providers to reduce data bandwidth, with a filter option (similar to a `WHERE` clause in an SQL statement)

9.3.4 Managing business performance with SAP and IBM Cognos TM1

The architecture shown in Figure 9-7 enables business analysts or planning controllers to perform planning scenarios on SAP Business Suite data with the market-leading IBM enterprise planning software, the TM1 toolset.

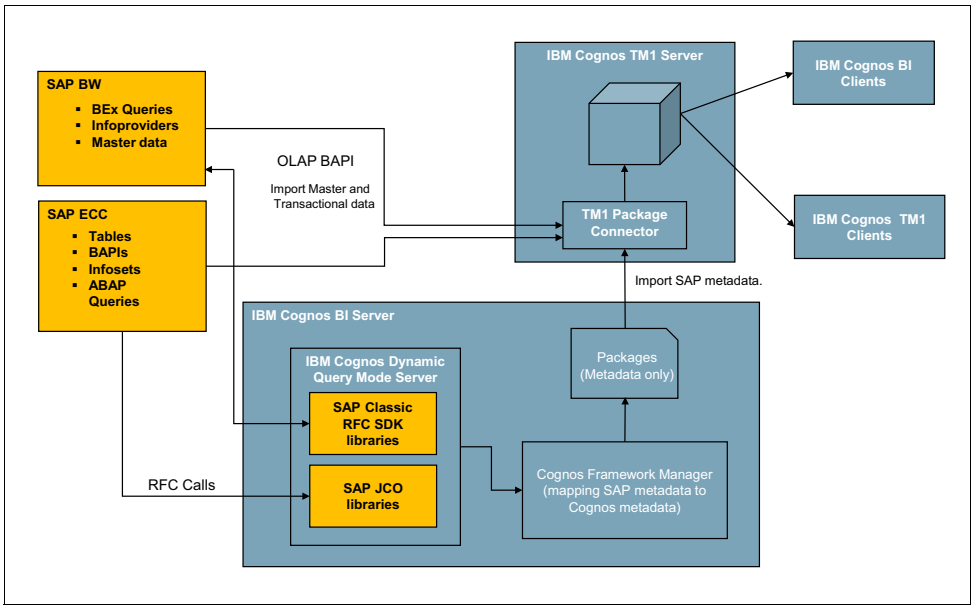

Figure 9-7 Managing business performance with SAP and TM1

TM1 uses the same SAP-certified interface used by the Cognos Business Intelligence platform to pull data into TM1 from SAP BW and SAP ECC quickly and efficiently. Using Cognos Business Intelligence packages with the IBM Cognos TM1 Package Connector, data is packaged and sent to TM1 using certified connections from SAP BW (OLAP BAPI) and SAP ECC (RFC).

Because the same OLAP BAPI interface is used to access data in SAP BW, there are no separate modules or programs to install on the SAP BW server. As a result, organizations gain the ability to use common structures in SAP BW, such as BEx Queries, InfoCubes, MultiProviders, Data Store Objects (DSOs), InfoSets, and Master Data objects.

For example, this integration architecture can be used as the premium approach if a business requirement exists to explore *what-if* scenarios with the TM1 planning toolset based on SAP BW aggregated budget or actual data with no major data manipulation upward. Cognos Dynamic Query server is not an ETL system, and therefore no complex rules can be applied between SAP BW and the Cognos server.

Compared to the streaming scenario for IBM Cognos reporting (described in 9.3.3, “Operational analytics with Cognos Business Intelligence directly accessing SAP solutions” on page 242), this integration architecture pulls data from SAP BW in a batch mode, which is a better fit for high data volume extractions.

This architecture uses the TM1 Package Connector ability to connect to an SAP ECC and SAP BW source systems through a published Cognos Business Intelligence package. The data is pulled from the SAP source system and persisted in TM1. The data is then consumed either through the TM1 toolset or through Cognos Business Intelligence reporting capabilities.

9.3.5 Predictive analytics with SAP

Predictive analytics is an advanced area of BA technology that produces a predictive score for each customer or other organizational element. Assigning these predictive scores is the job of a predictive model that, in turn, has been trained over the data. With predictive analytics, the enterprise learns from its cumulative experience (data), and takes action to apply what has been learned.

IBM SPSS Modeler is an extensive predictive analytics platform that is designed to bring predictive intelligence to decisions made by individuals, groups, systems, and the enterprise. By providing a range of advanced algorithms and techniques that include text analytics, entity analytics, decision management, and optimization, SPSS Modeler can help organizations to consistently make the correct decisions.

SPSS Modeler provides a visual, intuitive interface for predictive analytics, and it supports automated modeling and data preparation. It provides over 20 packaged predictive analytics techniques spanning classification, clustering, association, forecasting, and simulation. It uses natural language processing (NLP) to extract concepts and sentiments contained in text.

Entity analytics helps identify entities, for example customers, orders, employees, and so on, based on context, which provides richer insight into the entity itself while helping create more accurate models with cleaner data. SPSS Modeler provides the ability to extend analytical techniques using open source languages, such as R and Python:

R This open source statistical language is used by data scientists and expert analysts to create custom analysis routines and new algorithms.

Python This widely used, general-purpose programming language is often used for scientific and numeric computing.

SPSS Modeler delivers true enterprise reach, which enables you to access all enterprise data, structured and unstructured, from disparate sources. It provides a centralized, secure environment for managing and running models through IBM SPSS Collaboration and Deployment Services, and provides deployment features for integrating predictive analytics into business processes.

SPSS Modeler supports the SAP BW environment by providing a visual, data-independent data mining environment that can take advantage of both structured and unstructured data that is captured within SAP, or within other operational systems. SPSS Modeler connects to the SAP environment through an ODBC connection. Data is directly accessible, and can be analyzed, manipulated, and edited, assuming that the user has the appropriate credentials.

SPSS Modeler supports the ability to use SQL pushback to improve performance for large data sets. SQL pushback enables a user to push back key procedures. which could be data transformation or calculation of a predictive value through a model developed in SPSS Modeler. SQL pushback enables the system to run the commands directly on the data warehouse, rather than in-memory. This approach minimizes, and sometimes eliminates, data movement performed within SPSS Modeler, and improves performance significantly.

Models in SPSS Modeler Server enable the SQL statements to be generated, pushing back the model scoring stage to the database itself. For modeling streams that use these models, the full SQL of the stored procedure is pushed back to the database as SQL. The following list includes the models with these functions:

- C5.0
- Classification and regression tree (CART)
- Chi-squared Automatic Interaction Detector algorithm (CHAID)
- Quest
- Decision List
- Logistic Regression
- Neural Net
- Principal components analysis (PCA)
- Linear Regression

Testing has shown that most models improve performance up to 10 times.

For SAP HANA, SPSS Modeler is also able to push back models that include R code. With SPSS Modeler, users can run R syntax from R nodes in the user interface of SPSS Modeler. R model building can be run natively in the SPSS Modeler Server. The R syntax is parsed by SPSS Modeler and sent to the R program to process. R model scoring can either run natively in the SPSS Modeler Server (with the same technique used for model building) or with the R in-database scoring functionality.

For the R in-database scoring, R is present in the database to take advantage of the fact that processing can be done in the same database system that is storing the data. This capability is enabled for SAP HANA, and performance is greatly improved by reducing data movement.

Figure 9-8 shows the integration architecture for SPSS Modeler with SAP.

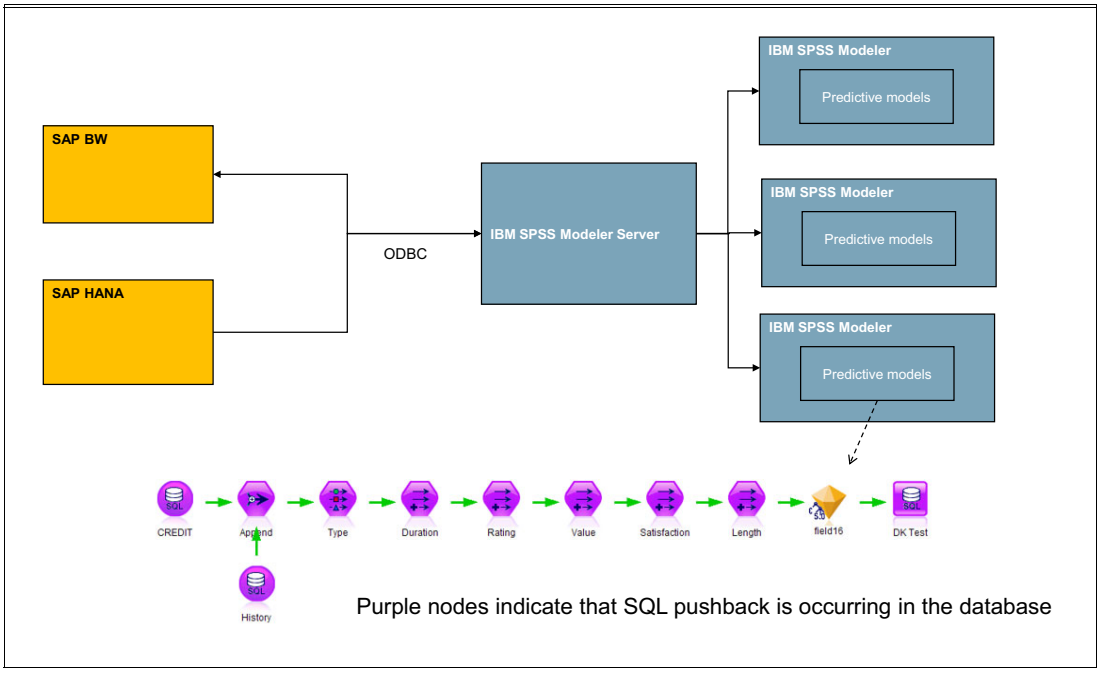

Figure 9-8 Integration architecture for IBM SPSS Modeler with SAP

This architecture uses ODBC connectivity with SAP, and illustrates how the SPSS ability to use the option of SQL pushback to the database greatly reduces network traffic.

9.4 Conclusion

IBM Business Analytics software provides established and mature solutions to enhance the value of large data volumes that are in SAP systems and applications, in addition to other heterogeneous data sources.

IBM Business Analytics solutions provide organizations the flexibility and agility to meet their many different business needs and requirements. This adaptability to the ever-changing needs of the business minimizes interruption to the SAP landscape, while accelerating time to deployment, and ultimately provides customers a faster return on investment (ROI).

IBM Business Analytics software delivers the actionable insights that decision-makers need to achieve better business performance. IBM offers a comprehensive, unified portfolio of BI, predictive and advanced analytics, financial performance and strategy management, governance, risk and compliance, and analytics applications.

With IBM software, companies can spot trends, patterns, and anomalies; compare *what if* scenarios; predict potential threats and opportunities; identify and manage key business risks; and plan, budget, and forecast resources. With these deep analytics capabilities, IBM customers around the world can better understand, anticipate, and shape business outcomes.

This chapter reviewed a set of integration scenarios that can help organizations choose the correct IBM integration framework certified by SAP. For direct-access options, the preferred approach is to select ABAP or BEx queries as SAP data providers, to efficiently reduce data bandwidth across the solution and avoid performance issues. In addition, it is also critical to perform a sizing exercise of IBM middleware architecture components based on the expected volume of data.

Note that the technical setup of SAP source systems for use with IBM Business Analytics software is nondisruptive, and does not require downtime. Shutting down the SAP server to implement the technical prerequisites is unnecessary, which is important when the production system is mission-critical and requires high availability (HA).

9.5 References

These websites are also relevant as further information sources:

- IBM Cognos Proven Practices: IBM Cookbook for IBM Cognos 10 for use with SAP NetWeaver Business Warehouse

  http://www.ibm.com/developerworks/data/library/cognos/infrastructure/cognos_specific/page551.html

- IBM InfoSphere Information Server

  http://www.ibm.com/software/data/integration/info_server/

- Predictive analytics on SAP with SPSS and InfoSphere Warehouse

  http://www.ibm.com/developerworks/data/library/techarticle/dm-1007predictiveanalyticssapspss/index.html?ca=dat

- IBM Cognos Mobile

  http://www.ibm.com/software/products/en/cognos-mobile

- IBM DB2 for Linux, UNIX, and Windows with BLU Acceleration speeds analytics

  http://www.ibm.com/software/data/db2/linux-unix-windows/db2-blu-acceleration/

- IBM InfoSphere Information Server Pack for SAP Applications

  http://www.ibm.com/software/products/en/infosphere-information-server-pack-sap-applications

- InfoSphere Information Server Pack for SAP BW

  http://www.ibm.com/software/products/en/infosphere-information-server-pack-sap-bw

- V9.1 IBM InfoSphere Information Server Integration Guide for Information Server Pack for SAP Applications (SC19-3876-00)

  http://www.ibm.com/e-business/linkweb/publications/servlet/pbi.wss?CTY=US&FNC=SRX&PBL=SC19-3876

10

DevOps for SAP

To realize the full value of an investment in SAP solutions, teams need to enable continuous delivery of change across their SAP landscape. The reason for this is simple: If an organization runs on SAP software, business success depends on how well delivery teams can manage change across the SAP landscape and the dependent non-SAP technology.

IBM Rational tools and technology enable you to achieve continuous delivery and reduce the cost and risk of managing change to the SAP landscape. IBM calls this set of tools and technologies *IBM DevOps for SAP*.

"Our collaboration with IBM Rational brings together the best of our combined application lifecycle management market leadership and can help customers reduce costs, manage change, and improve quality across the enterprise applications lifecycle." Dr. Uwe Hommel, SAP, Executive Vice President, Executive Board Member, Head of Active Global Support.[1]

This chapter describes how to achieve continuous delivery of new implementation and changes across the SAP landscape.

This chapter includes the following topics:

- 10.1, "IBM DevOps for SAP overview" on page 250
- 10.2, "Application lifecycle management for SAP" on page 253
- 10.3, "Collaborative development for SAP " on page 262
- 10.4, "Continuous testing for SAP" on page 266
- 10.5, "Continuous release and deployment for SAP" on page 268
- 10.6, "Continuous business planning for SAP" on page 269
- 10.7, "Summary " on page 272

[1] Source: *How do you address application lifecycle complexity?* http://www.ibm.com/software/rational/solutions/packagedapps/sap/

© Copyright IBM Corp. 2014, 2015. All rights reserved.

10.1 IBM DevOps for SAP overview

Most SAP projects involve some change to user and information technology (IT) processes, which can increase both cost and risk to the expected delivery of the business. The following list includes some of the top challenges that were reported by SAP customers about managing delivery of projects:

- Time and money spent on upgrades and updates. SAP applications typically require a steady stream of upgrades, updates, and bug fixes. Long change freeze periods and lengthy testing cycles can typically cause disruption to the business. Companies that fail to adequately test new releases risk exposing their business to potential application conflicts.
- Managing requirements and testing change. When companies introduce new technology, such as mobile apps, or upgrade their existing solutions, they must determine the potential effect on surrounding environments. Even minor alterations can require significant engagement with business stakeholders and partners, and involve considerable testing.
- Integrating with non-SAP technology. As non-SAP products, applications, and vendors are added to the IT landscape, companies must undertake extensive integration work and testing to ensure that all of the components function correctly. As the IT landscape grows in complexity, managing such cross-system integrations can be a difficult task.
- Collaborating with non-SAP teams. Where dependencies exist between SAP and non-SAP projects, SAP delivery teams must keep synchronized with the release cycles of other teams, enabling them to manage projects holistically. Without close collaboration and a shared toolset for change control, ad hoc IT change or application releases can cause business disruption.

By uniting the people, practices, technologies, and information that support the SAP implementation and change projects, the IBM DevOps solution helps project delivery teams promote a culture of continuous innovation and improvement, and can offer significant competitive advantage to their business:

- Cut costs. An IBM DevOps approach promotes a leaner approach to SAP development and delivery that reduces waste and rework, and enables project delivery teams to use existing resources more efficiently. By advocating extensive use of automation and virtualization capabilities, IBM DevOps for SAP makes it easy to set up new SAP development, test, and production environments for teams and partners, saving time and money.

  For example, in US dollars, $207 million of savings and $667 million in risk mitigation were made by an HR organization that used Rational software to optimize SAP delivery across their teams.
- Reduce risk. As the enterprise application environment becomes more complex, IBM DevOps can help to reduce the risk of business disruption. With a better view of the IT landscape, and of dependencies between processes and systems, teams can better manage the effect of IT change on the business. Comprehensive analysis and frequent integration testing helps to adopt a more proactive approach to change management, ensuring that new releases do not cause unexpected conflicts.

  For example, serious project overruns were addressed at a retail company that implemented Rational tools to address slipped deadlines and a perceived lack of quality in their implementation of SAP solutions. Both business and technical stakeholders improved confidence that their business can now operate as expected after SAP software upgrades or updates.

- Accelerate delivery. SAP teams are constantly being asked to shorten delivery cycles. With IBM DevOps, teams can apply a fast, iterative approach to upgrades, updates, integration, development, testing, and deployment that helps you keep up with changes to the business and the company's SAP environment.

  A faster time-to-value for process changes and releases helps ensure that businesses can readily take advantage of new SAP features and optimizations, providing businesses a fast-mover advantage to rapidly seize market opportunities and gain an edge on the competition. For example, the ability to determine the effect of change has improved from months to minutes for a large government agency that implemented Rational to identify and communicate how to deploy their SAP transformation.

By bringing development and operations teams closer together, and applying agile principles across the entire SAP software delivery lifecycle, IBM DevOps for SAP enables continuous innovation, feedback, and improvement for both business and IT stakeholders.

10.1.1 IBM DevOps for SAP key capabilities

As part of its approach to the DevOps concept, IBM outlines four key capabilities (steer, develop/test, deploy, and operate) that are required to support continuous delivery of SAP software. These capabilities are closely linked, with a continuous feedback loop, as depicted in Figure 10-1. This linkage and continuous feedback enable the business, development, and operations teams to learn and respond to user needs.

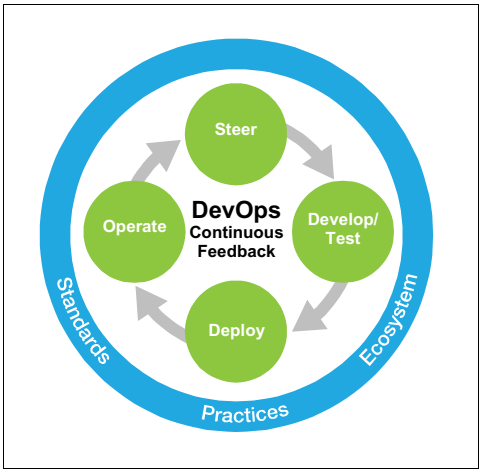

Figure 10-1 The IBM DevOps solution key capabilities

The following list describes these key capabilities:

- Steer

  Effective steering requires early and continuous feedback from users and the market. As the marketplace works today, and with the business climate we face, organizations need to ensure that they "do the correct things" and focus on activities where they will gain most value. Organizations capture business process changes, often in SAP Solution Manager, and define the business case or goals for IT delivery projects within the scope of the overall business strategy.

  When projects are run, changes and scope need to be managed in the context of the IT environment, and aligned with broader business goals. To achieve this, you need a big-picture view of your existing SAP and enterprise resource planning (ERP) applications. This will enable you to visualize and analyze your business requirements, and to plan and communicate any IT changes, upgrades, or deployments.

- Develop/Test

  Maintaining a competitive advantage requires the continuous innovation of ideas, and the ability to translate them into changes in the SAP landscape and dependent non-SAP technology. Collaborative development and continuous testing supports the evolution of a business idea into a high-quality SAP solution by applying lean principles, facilitating collaboration among all stakeholders, and striking the optimal balance between quality and time to market.

- Deploy

  Frequent delivery of SAP and dependent non-SAP systems requires automated, repeatable deployment processes. This improves SAP solution delivery and shortens time-to-value, which in turn drives down operating costs and reduces business risk. Automated deployments and middleware configurations can, if required, mature to a self-service model providing individual developers, teams, testers, and deployment managers with the ability to continuously build, provision, deploy, test, and promote SAP updates.

- Operate

  Operational, maintenance, and support teams need to understand the performance and availability of SAP applications and systems at all times. If a failure occurs or a bottleneck emerges, they must be able to identify the source of the problem, and take remedial action fast to avoid business disruption.

The key capabilities of DevOps for SAP are realized by the solutions that are shown in Figure 10-2 on page 253:

- Application lifecycle management for SAP. For more information, see 10.2, “Application lifecycle management for SAP” on page 253.
- Collaborative development for SAP. For more information, see 10.3, “Collaborative development for SAP ” on page 262.
- Continuous testing for SAP. For more information, see 10.4, “Continuous testing for SAP” on page 266.
- Continuous release and deployment for SAP. For more information, see 10.5, “Continuous release and deployment for SAP” on page 268.
- Continuous monitoring for SAP. For more information, see Chapter 12, “Systems management for SAP” on page 307.
- Continuous business planning for SAP. For more information, see 10.6, “Continuous business planning for SAP” on page 269.

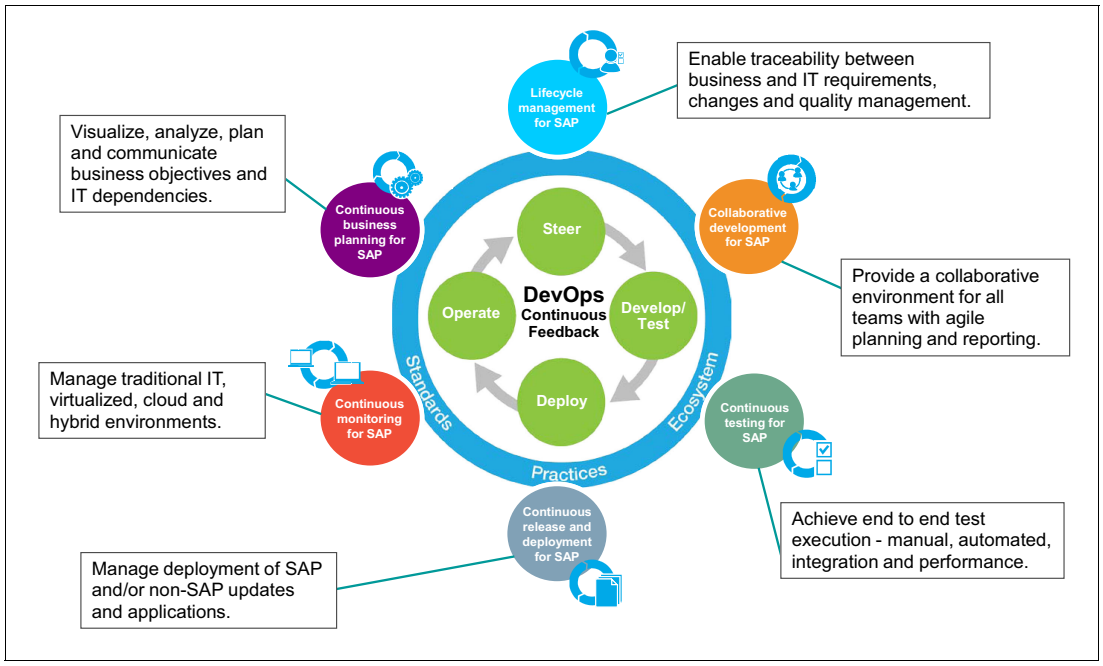

Figure 10-2 The solutions that achieve DevOps for SAP

10.2 Application lifecycle management for SAP

The goal of application lifecycle management (ALM) in the context of an SAP project is to provide consistent processes, methods, and tools for making communication friction-free. More than 80% of SAP implementations are heterogeneous in nature, and delivery teams require a single integrated solution across both the SAP and non-SAP components to achieve an end-to-end-centric transformation program.

Historically, ALM refers to an integrated process and tooling approach covering requirements management, change and configuration management, and test management. The approach of IBM to ALM is called *Collaborative Lifecycle Management* (CLM), because it emphasizes the need to have disparate teams interact and share artifacts and information throughout the delivery lifecycle.

CLM is also focused on the concept of *linked lifecycle data*. With linked lifecycle data, SAP artifacts have a unique Uniform Resource Locator (URL), and links are used to relate items to one another. This avoids traditional problems with data copying and mastership, and creates an open environment where data is separated from the tools and can be accessed using industry standard approaches, such as Representational State Transfer (REST).

CLM is defined by the following imperatives:

- Maximize product value with in-context collaboration.
- Accelerate time to delivery with real-time planning.
- Improve quality with lifecycle traceability.
- Achieve predictability with development intelligence (for example, rich and accurate reporting on development progress).
- Reduce cost with continuous improvement.

CLM is relevant in SAP business scenarios because, in IBM experience, enterprises have complex, heterogeneous middleware environments. Customers have multiple packages, multiple middleware, and multiple platforms that they are trying to integrate to fulfill various end-to-end business processes. The IBM CLM solution gives customers a single approach to managing SAP and non-SAP projects across their heterogeneous enterprise. Figure 10-3 shows the key components of IBM CLM.

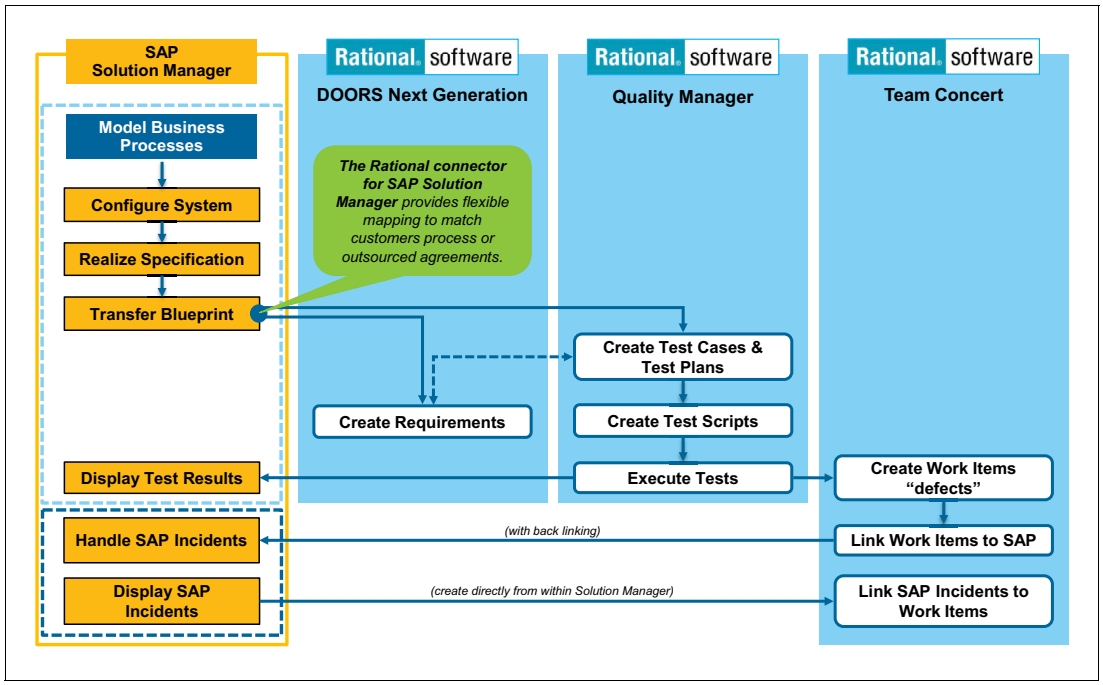

Figure 10-3 Extending SAP Solution Manager with IBM Collaborative Lifecycle Management

The following sections describe in more detail the IBM CLM for SAP software key solution capabilities, components, and products.

10.2.1 IBM Rational Connector for SAP Solution Manager

IBM Rational Connector for SAP Solution Manager can be used to integrate key areas of Rational software with the SAP ALM solution. Rational Connector for SAP Solution Manager enables the project delivery team to share data between Rational software tools and SAP Solution Manager.

Requirements managers and testers can export SAP Solution Manager blueprints to create test plans and test cases and requirements. Later, test results can be imported into SAP Solution Manager from IBM Rational Quality Manager for reporting and analysis. In addition, Rational Connector for SAP Solution Manager enables users to create links between IBM Rational Team Concert work items and SAP Service Desk incidents.

10.2.2 Requirements management

A comprehensive approach to defining and managing stakeholder needs, business requirements, gaps, and implementation solutions is needed for SAP and non-SAP systems. All of these forms of *requirements* should be formally documented, communicated, and tracked throughout the multiple phases of a typical delivery project.

The dynamic process definition and iterative blueprinting tools available in IBM Business Process Manager can substantially enhance the SAP process design approach, and complement the process design and SAP Solution Manager integration capabilities available in Rational tools.

IBM Business Process Manager enables you to easily run SAP processes without IT development, both at design-time and run time. This approach reduces SAP blueprinting time and risk, while positioning the SAP processes for the flexibility, agility, and control available with IBM Business Process Manager external process orchestration.

Traceability is paramount to ensure that business needs are matched by IT implementations. Being able to trace from business needs to business requirements, and all the way to test execution, ensures that the wanted changes are actually implemented in production systems. In addition, many enterprises have regulatory and compliance mandates that require audit trails documenting lifecycle traceability, and any changes made to a production system.

Within the context of SAP projects, *requirements management* addresses the following areas:

- Scope and business goals of the SAP development project
- Business requirements and business rules
- SAP requirement gaps

Collectively, these activities enable the project to capture an enterprise-level business blueprint that is traceable through project planning, execution, and all forms of test. Figure 10-4 shows an example of how team leaders can gain visibility into coverage and completeness of a project plan.

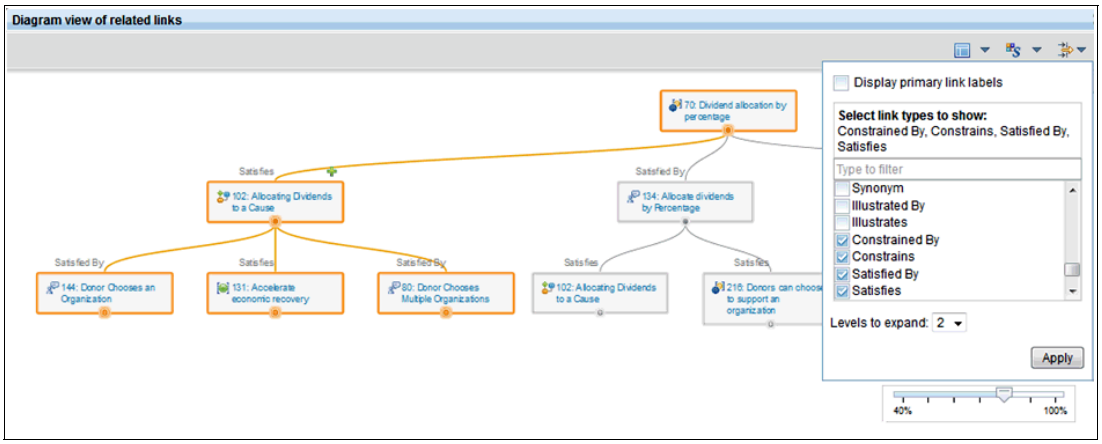

Figure 10-4 Requirements traceability example

This integrated solution proactively responds to gaps (highlighted in different colors) as they surface throughout the project. Such issues can be quickly identified and resolved.

The SAP business blueprint contains business requirements and IBM Rational DOORS® Next Generation can effectively manage any form of SAP business requirements definition. IBM fully supports the SAP-mandated business requirements management process (blueprinting) based on SAP Solution Manager.

The main scenarios when SAP Solution Manager is integrated with Rational CLM tools are as follows:

- SAP Solution Manager is used to manage all SAP-related blueprint items. DOORS Next Generation is used to manage requirements for all requirements related to non-SAP components of the overall solution.
- SAP Solution Manager is used to document the structure of an SAP business process hierarchy. However, all content in the business process is managed in DOORS Next Generation as structured data. DOORS Next Generation provides further structure and decomposition to the business process hierarchy managed in SAP Solution Manager, for example, more detailed requirements decomposition.

  This approach adds value by providing more structure, traceability, and integrated management of actual contents of the business process requirements when compared to using SAP Solution Manager alone, where requirement details are managed manually as basic document attachments.

The preferred approach is to use SAP Solution Manager to capture business process hierarchy for SAP blueprint, and to use DOORS Next Generation to elicit, communicate, and manage actual requirements as structured data.

DOORS Next Generation enables teams to define, manage, and report on requirements throughout the project lifecycle. The defined requirements can be traced through the various levels of the requirements matrix and across the involved components. DOORS Next Generation integrates with SAP Solution Manager to link artifacts between the SAP Solution Manager and IBM tools.

For SAP projects, the requirements hierarchy maps to the elements of the SAP Blueprint phase: process requirements, Workflow, Report, Interface, Conversion, Enhancement, and Form (WRICEF) requirements, component requirements, and so on. Process requirements are defined at a high level and detailed through business process models. WRICEF requirements identify gaps in the existing process. Each WRICEF is further detailed in component requirements and other functional and non-functional requirements.

10.2.3 Blueprint push from SAP Solution Manager

Through IBM Rational Connector for SAP Solution Manager (jointly developed with SAP engineering), project teams can link elements in the SAP Solution Manager business blueprint to elements managed by DOORS Next Generation and Rational Quality Manager.

SAP guidelines are to have the business blueprint defined and accessible from within SAP Solution Manager. The use of DOORS Next Generation does not contradict these guidelines. The business process hierarchy can be defined in SAP Solution Manager and mirrored (using IBM Rational Connector for SAP Solution Manager) in DOORS Next Generation. Doing so enables a consistent approach for putting the SAP business blueprint within the overall context of enterprise business goals, requirements, tests, and processes.

In Figure 10-5, SAP blueprint is mirrored in DOORS Next Generation using IBM Rational Connector for SAP Solution Manager. This process is often referred to as a *Blueprint push*:

- IBM Rational Connector for SAP Solution Manager automatically creates requirements, test plans, and test cases in the CLM project.
- All data is linked for traceability using the Open Services for Lifecycle Collaboration (OSLC) standard.
- Requirement collections are used to structure the requirements (business scenarios, processes, and steps).

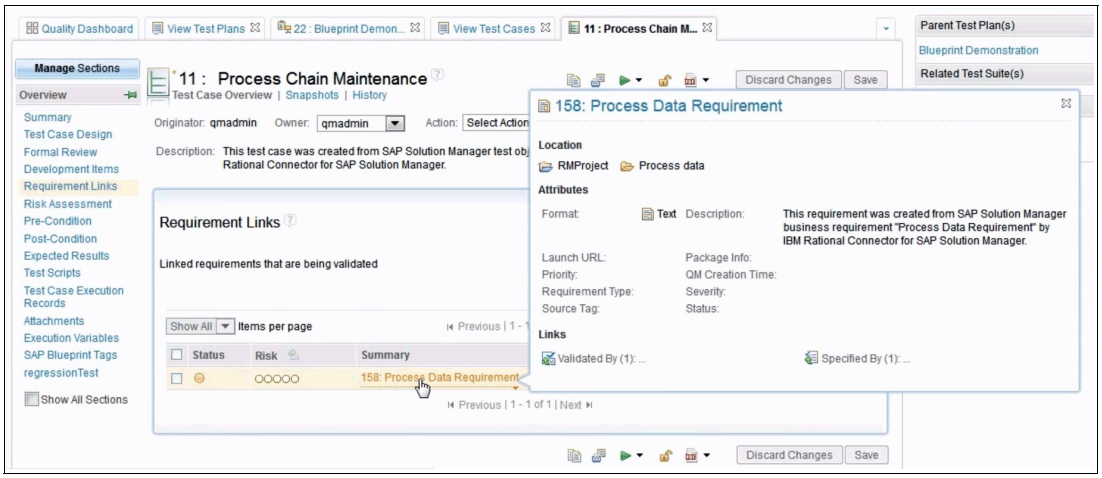

Figure 10-5 Results of "Blueprint push" from SAP Solution Manager to DOORS Next Generation

10.2.4 Project planning and execution

Because customer environments are inherently heterogeneous, SAP delivery projects often involve work affecting multiple disparate applications and systems in addition to the core business processes of the enterprise. With the IBM CLM solution for project planning and execution, teams can use a single tool for project and iterations plans, work item tracking, deliverable activity management, and release planning.

Rational Team Concert covers multiple areas of functionality common in all variations of ALM:

- Work item and deliverable activity management
- Project planning
- Source control management
- Continuous build and integration
- Dashboards and reporting
- Method and process enforcement

Figure 10-6 shows how Rational Team Concert can be used as the collaborative project hub to track all work, control project governance, and identify gaps.

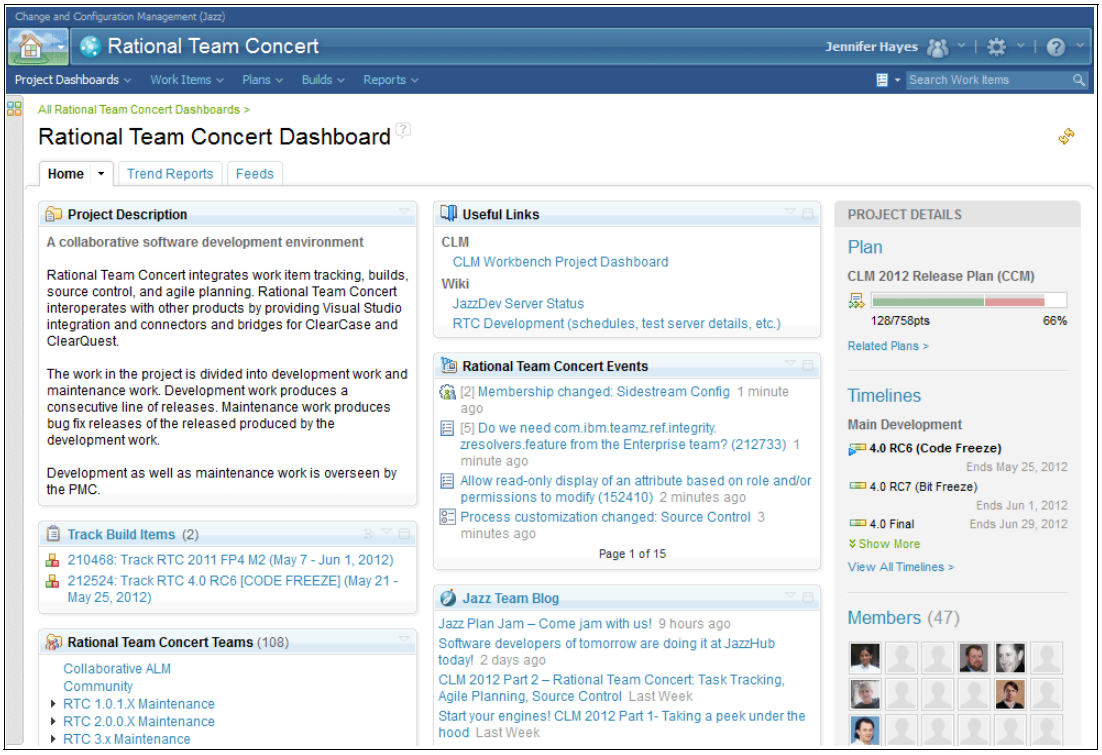

Figure 10-6 Rational Team Concert: Single tool, many capabilities

Rational Team Concert can be used to manage SAP and non-SAP projects in a unified way. This approach enables teams to plan and run projects based on end-to-end business processes, and to coordinate all changes and release deliveries across the different applications and systems. Rational Team Concert is also used as the collaborative project hub to track all work, control project governance, and identify gaps in needed work items for completion. Figure 10-7 shows how Rational Team Concert can highlight planning gaps.

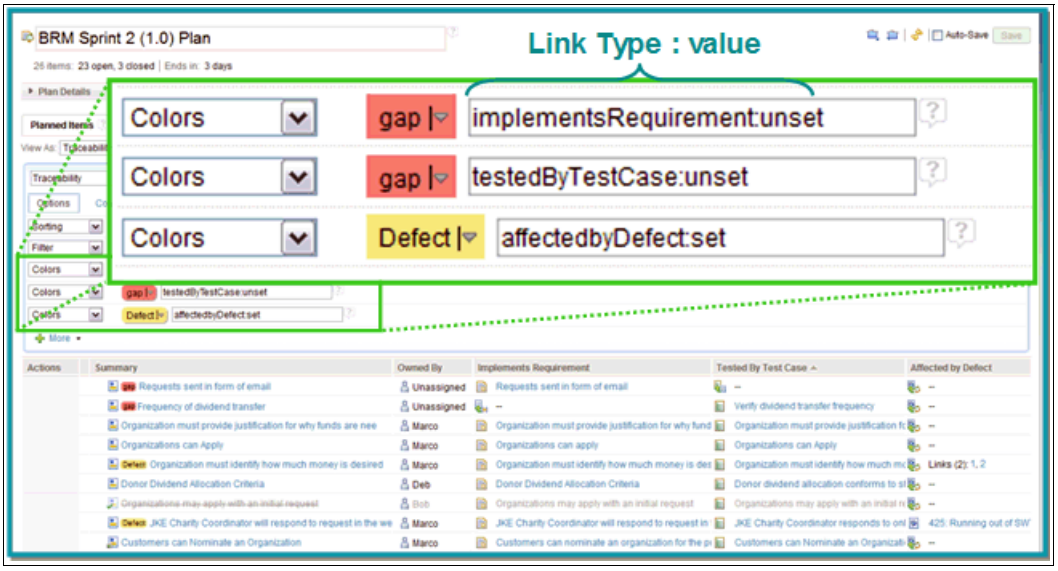

Figure 10-7 Tailoring traceability view in Rational Team Concert

10.2.5 Change, defect, and incident management

Enterprise integration projects go through several phases, such as definition, development, testing, delivery, and so on. Team members collaborate to ensure that all changes and implementation dependencies are handled correctly. Change management is a key component in governance. Effective *change request management* requires a common information repository that all team members can access.

Change requests have a governance lifecycle that is specified by the approval process for that type of request. Associated with the change request are all of the artifacts related to that request, such as requirements, code, models, documents, and so on. Trace matrixes show the effect that a change has on other project artifacts.

Rational Team Concert is used to define and govern changes throughout the lifecycle. Rational Team Concert provides a single change management platform for both SAP and non-SAP change requests. Change requests managed by Rational Team Concert are not limited to source code changes, but enable you to manage any changes in SAP software, for example, user interface (UI) changes. To facilitate change governance process, lifecycle management workflows can be attached to change requests.

Figure 10-8 shows how Rational Team Concert can be used to identify and select appropriate delivery processes. Rational Team Concert provides multiple *process templates* ready-to-use that can be used as a starting point. Process templates provide a blueprint for the initial process of a project area and iteration structure.

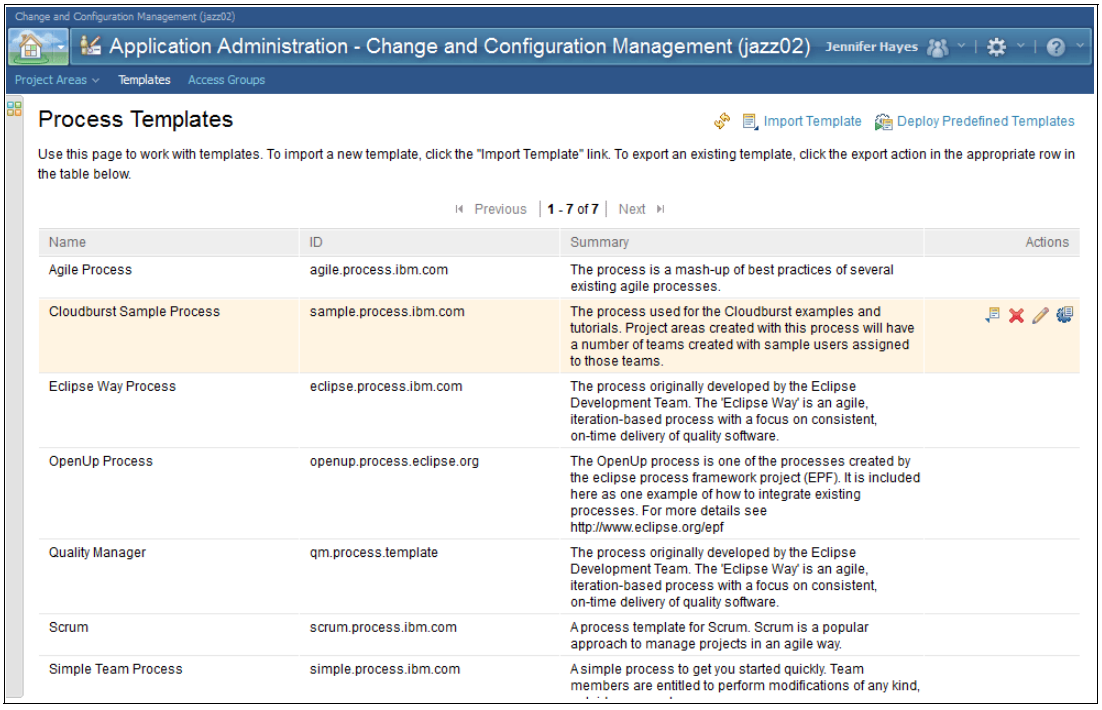

Figure 10-8 Lifecycle process in IBM Rational Team Concert

In addition to change governance, changes need to be delivered to their representative applications and systems. Rational Team Concert is used to deliver changes and manage the version control of related artifact changes, such as source code. For example, consider a particular lifecycle iteration (or even phase) focused on the update of a critical business process, such as order-to-cash (O2C).

In this example, O2C touches multiple systems, applications, and packages. Rational Team Concert project plans contain the work items, stories, use cases, and tasks associated with that iteration of the O2C update. Each work item can contain links to the associated other artifacts that it is related to, such as test plans, test execution results, requirements, business goals, and so on. Software changes (change sets) can also be associated to specific work as can all collaboration and discussions among team members.

Using Rational Team Concert's build capability (along with other associated deployment tools), work can be delivered in a synchronized fashion to all of the SAP and non-SAP applications and middleware systems affected by the O2C upgrades.

Finally, *defect management* is typically considered a variation of change request management. Rational Team Concert supports an integration with the SAP Service Desk, so defects and enhancements that need to be coordinated with the SAP help desk can be automatically managed and tracked. For example, Figure 10-9 shows a specific defect that is forwarded to the SAP Service Desk. The defect submission form is populated with live data from SAP Service Desk.

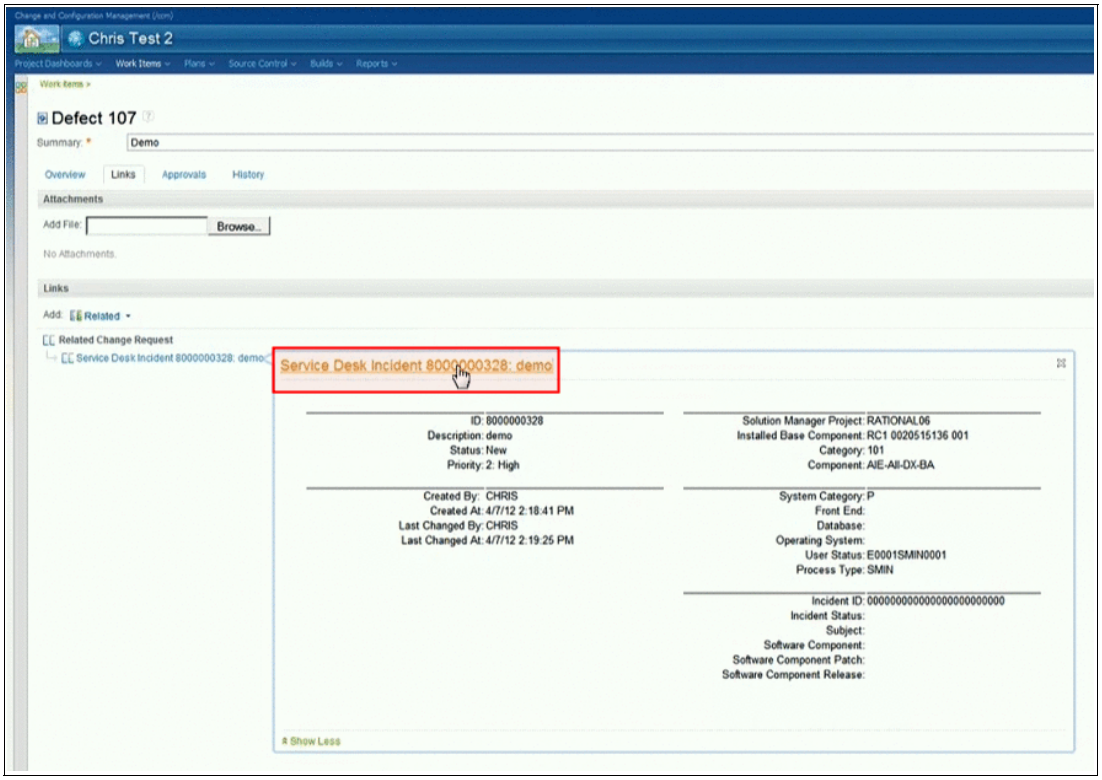

Figure 10-9 IBM Rational Team Concert and SAP Solution Manager Service Desk

10.2.6 Quality management

Quality management consists of the following main functional activities:

- Test planning
- Test execution
- Test reporting

In addition, defect management is typically viewed as an extension to quality management although within the context of ALM it is usually captured as a specific variation of change request management.

With test planning, you create your test plans and test cases. Test cases are turned into configured test cases when they are attached to a test script. This test script can be a manual test procedure, or an automated functional or performance test. Testers can then group configured test cases into test suites for execution. Alternatively, test cases can be run individually.

Remember: Rational Quality Manager is one of the testing and quality solutions referenced by SAP that provide extended capabilities beyond SAP's own solution built into SAP Solution Manager.

"By standardizing your testing program and encouraging collaboration between business stakeholders and the IT team, SAP Solution Manager and IBM Rational software help to minimize errors and other risks while maximizing business results. By identifying and managing risk more effectively, you can make better decisions about your testing priorities. Providing end-to-end traceability of test results back to the business requirements lets you confirm that the requirements have been addressed and implemented properly without adversely affecting other areas of the software." [2]

Important: Rational Quality Manager is used for SAP *and* non-SAP-centric quality management, and specifically, for test planning and reporting.

Rational Quality Manager is fully integrated with the IBM concept of open linked data, and has a transparent integration to both Rational Team Concert and DOORS Next Generation. The implication is that data managed in any of these tools can transparently link to, and be reported on across, all other repositories.

IBM Rational Connector for SAP Solution Manager provides an integration between SAP Solution Manager and Rational Quality Manager. Objects in the SAP business blueprint are mapped to test plans and test cases, and test results are automatically synchronized back into SAP Solution Manager at the appropriate level within the blueprint. This approach enables the Blueprint to be a general business-focused container for the overall test architecture, which is often referred to as a Blueprint push.

10.2.7 Impact analysis

Rational Quality Manager supports an integration into both SAP Business Process Change Analyzer (SAP BPCA) and IBM InfoSphere Optim™ System Analyzer for SAP Applications.

SAP BPCA enables customers to compare multiple versions of an SAP application and system, and generate a change set related to specific business processes defined in the SAP business blueprint. This set can be automatically related to test plans and test cases defined in Rational Quality Manager.

[2] Source: *Teaming SAP Solution Manager and IBM Rational Software for Top Test Management*

Figure 10-10 shows how SAP BPCA and Rational Quality Manager can be used to identify the minimum amount of testing required after a change to the business process hierarchy.

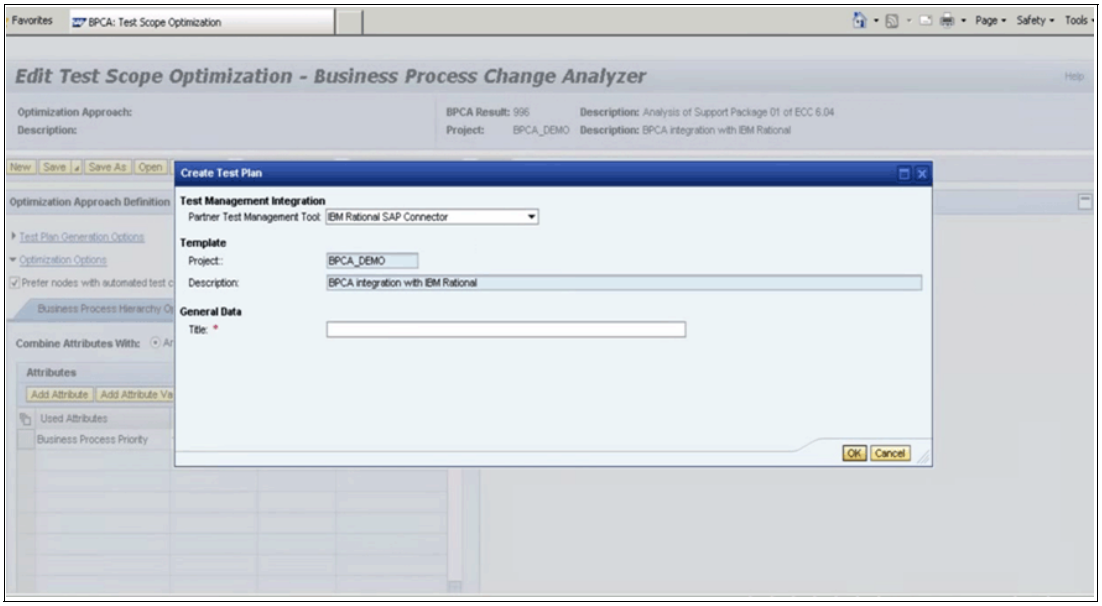

Figure 10-10 Testing recommendations based on changes to the business blueprint

10.3 Collaborative development for SAP

Agile software development and delivery is mainstream for custom applications, and is increasingly being used to successfully develop and deliver SAP customizations, integrations, and custom development. The collaborative development for SAP solution from IBM Rational brings agile development to both SAP and non-SAP development and delivery.

The essence of agile software development is to accommodate change by more quickly developing and delivering SAP customization, integrations, and custom development in short time scales. Because agile development is not prescriptive (*do this, and then do that*) a variety of methodologies and supporting techniques have emerged, including agile variants of SAP AcceleratedSAP (ASAP) methodology to guide and enable SAP teams.

Agile methods, such as Scrum, use *user stories* to specify what needs to be delivered. These user stories accumulate in a non-ordered list called a *backlog*. Backlog activities, such as *ranking* user stories with respect to each other and complexity estimation (*story points*), are continually done to plan and deliver a release. Releases have timelines, and are delivered by projects involving teams of *generalized specialists* in individual roles.

In the Scrum method of agile software development, work is confined to a regular, repeatable work cycle, known as a *sprint* or iteration. A sprint is a short period of time (typically 2 - 6 weeks) where planning, development, testing, and delivery of running software takes place. A release typically consists of many successive sprints. The objectives for a sprint can vary, often depending upon its position in the release lifecycle or *timeline*. Agile project progress is measured by delivery of *running and tested* software and related work.

Agile development has proven benefits, successes, and a focus on accommodating change throughout the lifecycle, and on quick delivery of working, valuable software. Therefore, the reasons why organizations, teams, and individuals are drawn to the agile approach, and eager to apply it to the development and delivery of SAP software can be easy to understand.

Figure 10-11 shows how agile processes are managed and supported in Rational Team Concert.

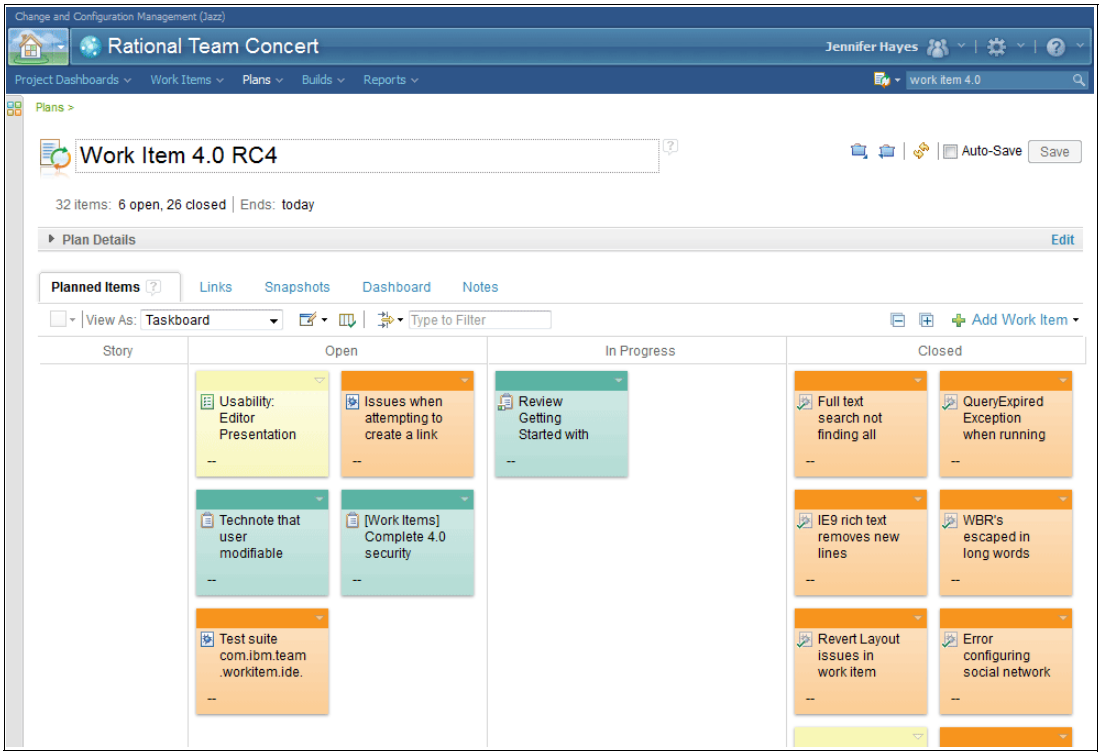

Figure 10-11 Agility at scale for SAP projects

The collaborative development for SAP solution provides immediate value for organizations that require development projects to extend the use of their SAP applications. Collaborative development for SAP solutions from IBM includes ready-to-use, customizable agile and lean planning, change, work, delivery management, and execution methodology support, with content and capabilities that provide the following benefits:

- Improve SAP developer and team productivity using integrations with SAP eclipse-based High-Performance Analytic Appliance (HANA), Advanced Business Application Programming (ABAP), and NetWeaver integrated development environments (IDEs) for improved developer productivity.
- Enable better decisions based on real-time, transparent visibility into SAP Agile delivery and maintenance projects.
- Accelerate agile adoption and results using pre-configured and customizable agile (or other) method definition and automated enactment.

The core product of collaborative development for SAP solutions is Rational Team Concert, a market-leading agile project planning, change, defect, and delivery management solution.

Figure 10-12 shows that Rational Team Concert provides a single, consistent project planning, execution, and change management solution that integrates (shell-sharing) with the SAP Eclipse-based IDE.

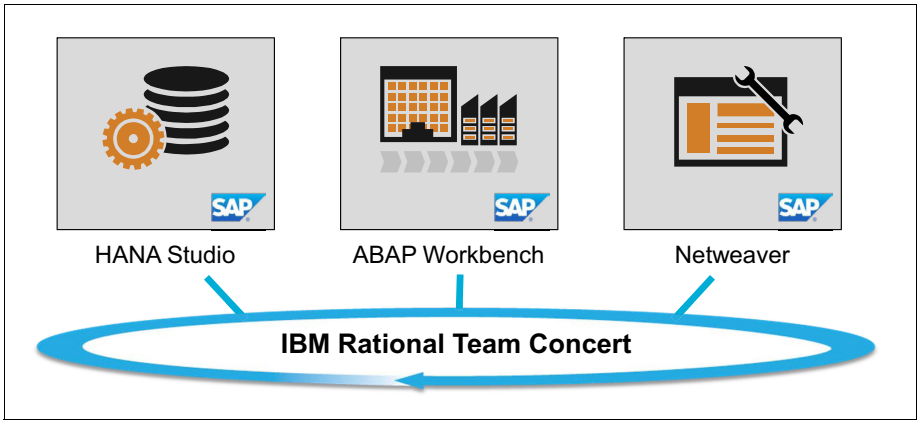

Figure 10-12 In-context agile planning and change management with Rational Team Concert integrated with SAP Eclipse IDE

The following sections describe in more detail the key solution capabilities, components, and products of the IBM Rational Collaborative development for SAP solution.

10.3.1 Improve SAP developer and team productivity

By integrating agile project work assignments (for example, user stories and related tasks, defects, and so on) and methods directly into the SAP Eclipse-based IDEs, the IBM collaborative development for SAP solution enables SAP developers to collaborate with others, and to manage their work assignments and deliverables in a single environment that improves developer productivity and results. Agile project and team queries, dashboards, and reusable assets are also available from this same environment.

Figure 10-13 on page 265 shows an example of Rational Team Concert being used within SAP NetWeaver Developer Studio.

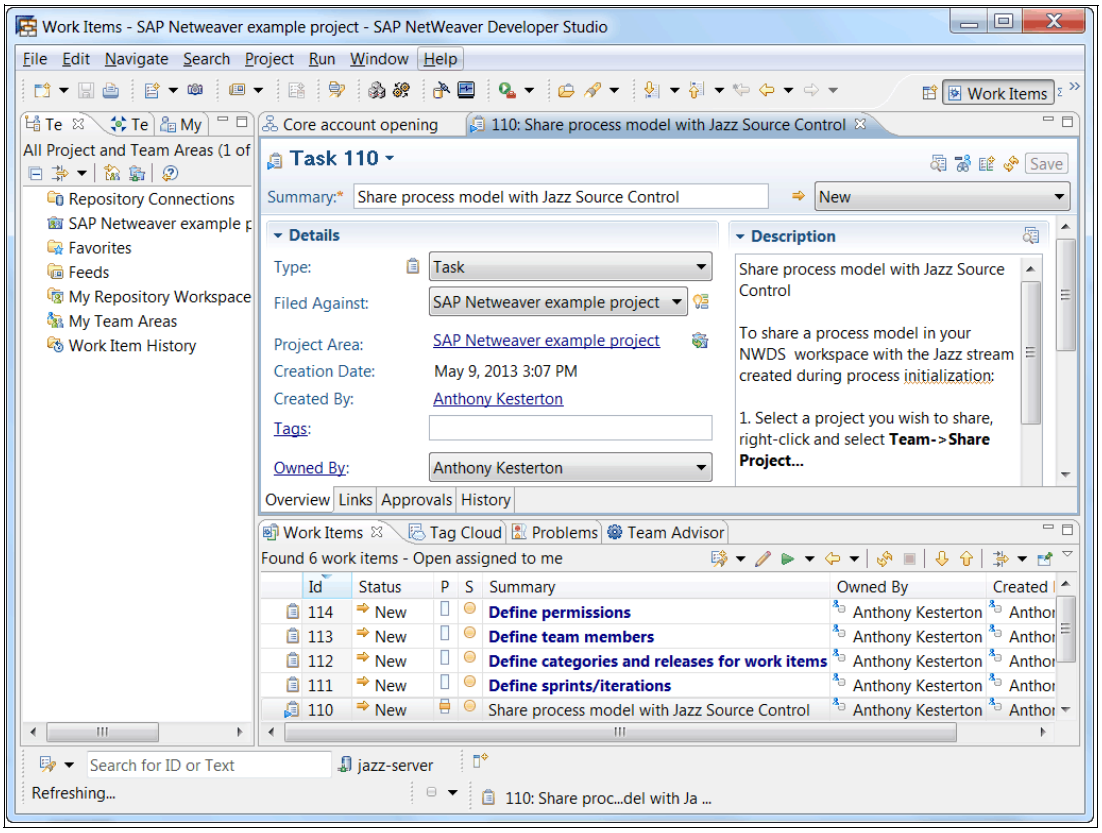

Figure 10-13 Rational Team Concert and SAP NetWeaver Developer Studio

SAP (and non-SAP) developers, testers, other team members, and stakeholders can use their web browser to use the IBM Collaborative development for SAP solution.

As SAP developers and teams go about doing their work using the IBM Collaborative development for SAP solution, it automatically captures and persists information about work assignment status, project plan progress, and other related measures and metrics in common, shared operational and analytical repositories.

10.3.2 Real-time visibility into SAP Agile delivery and maintenance projects

Business and technical stakeholders can benefit from real-time visibility on the status of the project. Without slowing down development teams with requests for information and reports, Rational Team Concert can be used to perform the following tasks:

- Get answers to project-relevant questions:
  - What SAP configuration and custom development work is in scope for this release or for this sprint?
  - What work has been approved? Completed? Tested?
  - What is the project status with respect to the planned and approved timeline and scope?

  This information is accessible from web-based dashboards with predefined and customizable reports (graphs, lists, and pie charts) that can be drilled down into for further detailed information.

- Have better collaboration, visibility, and control over SAP and related resources and work done across internal or external organizational (business unit, system integrator (SI), independent software vendor (ISV), and so on), geopolitical, and other types of boundaries. For example, determine what a specific SI is responsible for, and what the SI's status is.
- More effectively run and scale SAP Agile delivery and maintenance projects by using ready-to-use support for Kanban (a scheduling system for lean and just-in-time (JIT) production), other lean practices, and other proven agile scaling techniques (such as the ability to define multiple backlogs, customized SAP delivery roles, and so on).

10.3.3 Accelerate agile development adoption and results

Agile delivery of SAP development and testing is emerging as a way to address the need for accelerated and continuous feedback with the business. Updates, fixes, process changes, extensions, integrations, and customizations can be provided on a continuous basis for both SAP and non-SAP systems.

This approach increases the need for earlier feedback on delivery progress. Continuous delivery requires greater maturity in the testing process, which requires the introduction of continuous automated and manual performance, integration, and mobile device testing. Uniting information from across the whole SAP development, test, and deployment lifecycle begins to support improved collaboration.

10.4 Continuous testing for SAP

Test execution refers to running manual, functional, integration, mobile, or performance test scripts. The basic hierarchy of SAP test information is that a test plan contains multiple test cases, which contain multiple test scripts. Test execution documents the running (manual or automated) of a test script. The test execution record results are linked to the test case, along with any associated defects created as a result of a test failing.

"Combining the SAP Solution Manager application management solution with IBM Rational software results in comprehensive, automated, and integrated support for your test management program. This approach can help to minimize errors, reduce costs, and align your business and IT goals more effectively."[3]

The following sections describe in more detail the continuous testing for SAP software key solution capabilities, components, and products.

10.4.1 Functional testing

Functional testing refers to user-focused testing. Within an SAP context, this is usually driven by the business processes defined in the SAP Business Blueprint. Rational Quality Manager has built-in manual testing that, combined with Worksoft Certify, enables testing of end-to-end SAP business processes, objects, transactions, and screens, including UI5. It uses a ready-to-use testing framework for many SAP applications, eliminating the need, in most cases, for complex programming and test script capture.

This approach enables the creation of regression test suites that span multiple SAP release deliveries of the same business processes.

[3] Source: SAP Solution Brief: *Teaming SAP Solution Manager and IBM Rational Software for Top Test Management*

10.4.2 Integration testing and service virtualization

Integration testing refers to the ability to virtualize and simulate the interaction between systems and applications. IBM Rational Integration Tester, part of IBM Rational Test Workbench, includes comprehensive support for the following protocols:

- SAP Process Integration (PI)
- SAP Business Application Programming Interface (BAPI)
- SAP Intermediate Document (IDoc)
- SAP Remote Function Call (RFC)
- More than 100 other non-SAP messaging protocols

This enables for the virtualization of complex system connections when specific applications might not be available or stable yet in the release cycle.

IBM Rational Test Virtualization Server enables the deployment of virtualized services, software, and applications for simplified, efficient testing. It accelerates the delivery of complex test environments, and enables complete integration testing earlier and more frequently in the development cycle.

Figure 10-14 shows three common scenarios for integration and virtualization testing of SAP landscapes.

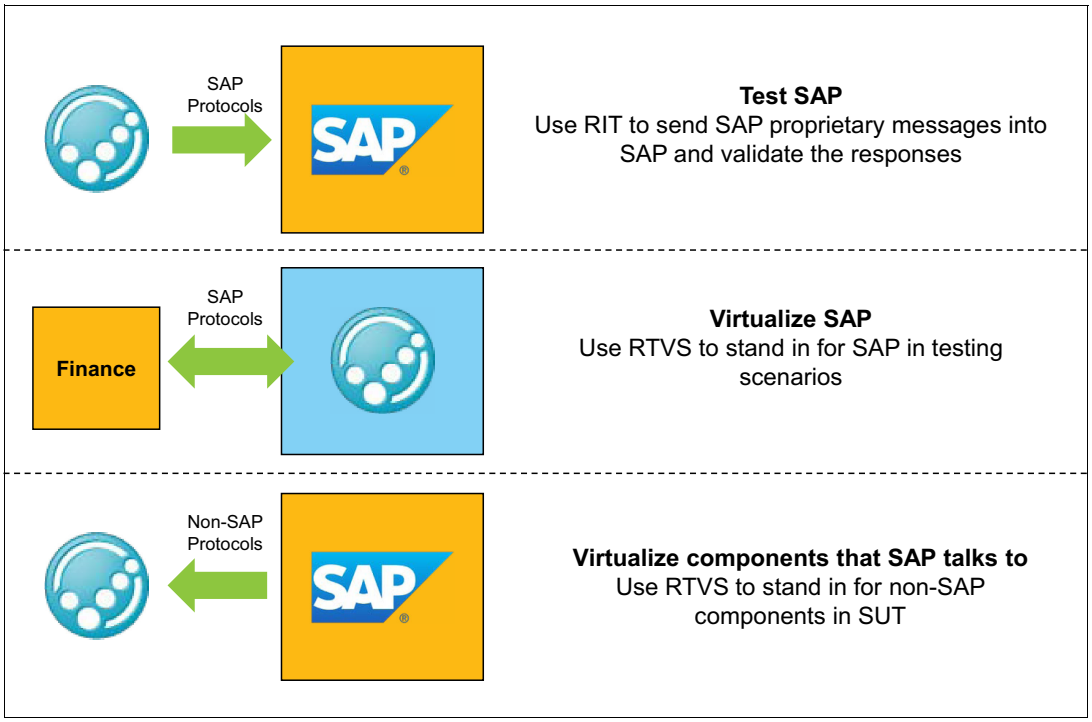

Figure 10-14 Integration and virtualization scenarios for SAP landscapes

10.4.3 Performance testing

Performance testing refers to being able to measure application and system response to load and stress. Rational Performance Tester has comprehensive support for SAP and non-SAP systems, and can provide insight into database, communication protocol, and server response time delays.

"Rational Performance Tester is one of the best performance testing tools I have ever come across. The dedication and expertise of the IBM Rational support team was outstanding." Dr. Peter Jäger, Senior Technology Consultant, Value Prototyping Performance & Benchmark Labs, SAP[4].

10.5 Continuous release and deployment for SAP

Often the most problematic stage of any SAP project is the deployment and release process. The phrase "it works for me" is the worst possible excuse for the failure of both development and operations teams to deliver on adding important improvements to the business. When businesses run on SAP, their success depends on how well they manage IT change. An inability to continuously deploy can negate the investment they make in developing and implementing change into business processes and integrations.

Costly, error-prone manual and duplicative processes delay innovation and affect competitiveness. Slow deployment to development and test environments leaves teams waiting, lowering productivity and affecting the entire business. Implementation teams need to deploy updates, changes, integrations, and fixes continuously to both SAP and non-SAP environments, and to do it all with limited resources and within ever-shrinking time frames.

Traditional approaches to SAP deployment rely largely on manual processes and spreadsheet management. These approaches can be error-prone and time-consuming, delaying the response to business needs and IT change requests.

Silos of processes, projects, people, and tools can make it difficult to gain a clear view of what needs to be deployed, and to be confident that the wanted objective will be achieved. Organizations looking to improve the value delivered by SAP solutions need to improve release management, reduce deployment times, streamline the deployment process, and reduce production outages.

SAP application landscapes are most often composed of many different applications that are interdependent, and can often involve both SAP and non-SAP applications and components. The components might be separate in that the IT organization can have different development teams working on each of them, but they are also interrelated in that no single application component by itself is a fully functioning solution. In addition, each SAP or non-SAP component can be released at a different time, and can be on different versions.

A flexible deployment and release strategy is to define for each component in the overall IT (and SAP) landscape its own deployment process. This approach enables the application component deployment processes to be orchestrated in different, reusable combinations that define an integrated release, so that an entire solution can be deployed and released together. This flexibility enables IT organizations to create deployment processes that they can reuse every time they deploy.

[4] Source: *Load testing SAP ABAP Web Dynpro applications with IBM Rational Performance Tester* at http://ibm.co/1oi3dZZ

Therefore, IT organizations can rehearse the deployment processes and develop a high degree of confidence in them. They can take a similar approach to defining and reusing target deployment environments. The result is a managed inventory of highly flexible, reusable combinations of component and application deployment processes and target environments.

The IBM Rational Continuous Deployment and Release for SAP solution can be used to coordinate the deployment and release of SAP and non-SAP landscapes. Specifically, it provides capabilities to perform the following tasks:

- Plan, coordinate, orchestrate, and manage releases of integrated application deployments.
- Automate the deployment of required non-SAP configured middleware.

For example, consider that a project team might be deploying an SAP enhancement package for SAP customer relationship management (CRM) for SAP HANA, where SAP HANA consists of the required hardware and the SAP HANA appliance software. In this example, you expect to replace the underlying relational database of the SAP CRM system with SAP HANA. When you deploy, you do not need to redeploy the web server component, only the database component.

IBM UrbanCode™ Deploy can inventory environments and identify if a supported web server component is at the correct version. Later, you might deploy the application to a pristine quality assurance (QA) environment. Again, IBM UrbanCode Deploy has the ability to inventory the environment, and to realize that changes to both the database and the web server components need to be deployed. IBM UrbanCode Deploy will therefore deploy them to both components in the QA environment.

The results are quicker releases, improved communications, and less risk associated with your releases.

10.6 Continuous business planning for SAP

Large investments by SAP customers are often made to enable a vital business transformation, which simply has to work. The challenge is that many ways of deploying IT are often available to successfully enable such a transformation.

A properly executed business planning initiative can provide the “big picture” and framework to select the best overall route. With continuous business planning, organizations can ensure that the SAP business model as defined is the business model that they implement, and can guide SAP investments to execute outcomes that support their corporate strategy.

This solution enables organizations to accomplish the following objectives:

- Assemble and organize business functional modules into an enterprise system where dependencies are captured between SAP processes and supporting data, which is also directly imported from SAP ERP systems.
- Get the design correct, and validate the business blueprints before implementation.
- Visualize, analyze, plan, and communicate SAP deployment on an ongoing basis, to ensure that the solution maintains its wanted outcomes over time.

The following sections provide more detail about the key solution capabilities, components, and products of the continuous business planning for SAP solution.

10.6.1 Enterprise architecture

IBM Rational System Architect is an enterprise architecture tool that enables organizations to analyze and plan their business and SAP technology architecture. Figure 10-15 shows the typical data points used to visualize the enterprise architecture.

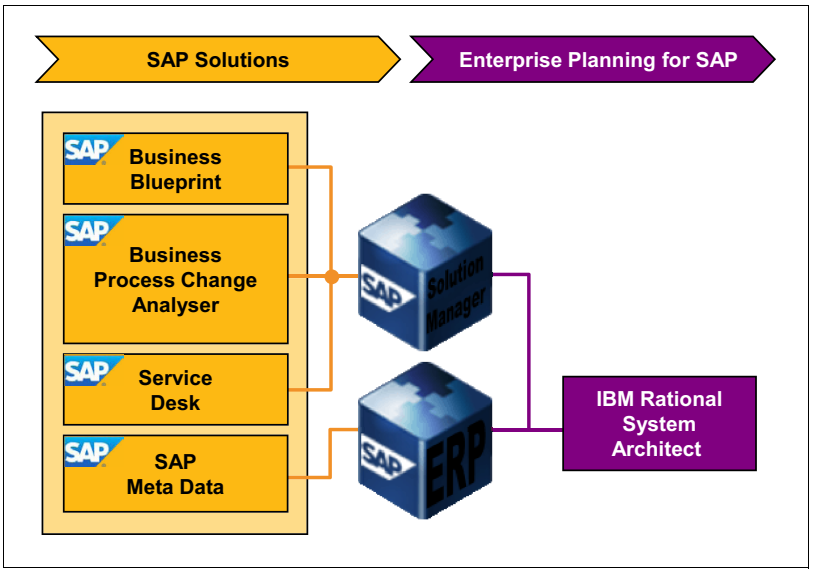

Figure 10-15 Continuous business planning for SAP

Rational System Architect helps organizations understand how their SAP systems map to their overall architecture (see Figure 10-16), what happens when architectural changes are made, and how the SAP environment can be used across the enterprise.

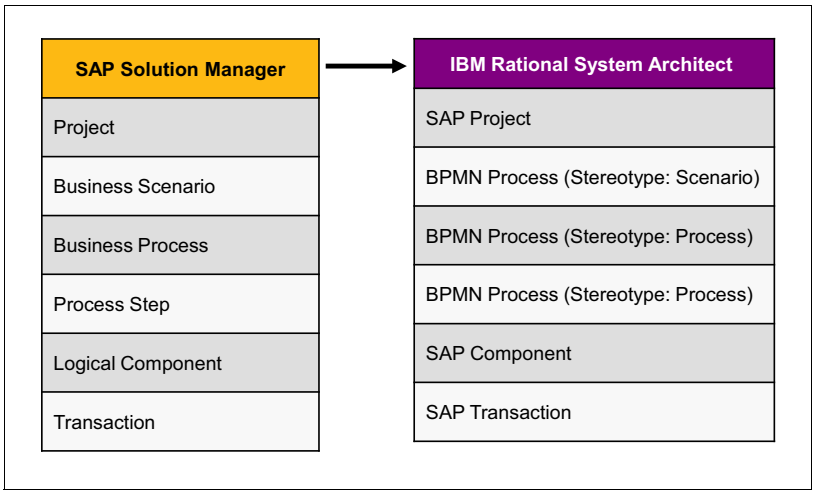

Figure 10-16 Mapping the SAP Solution Manager architecture in Rational System Architect

Figure 10-17 shows how key components of the SAP Solution Manager architecture are connected.

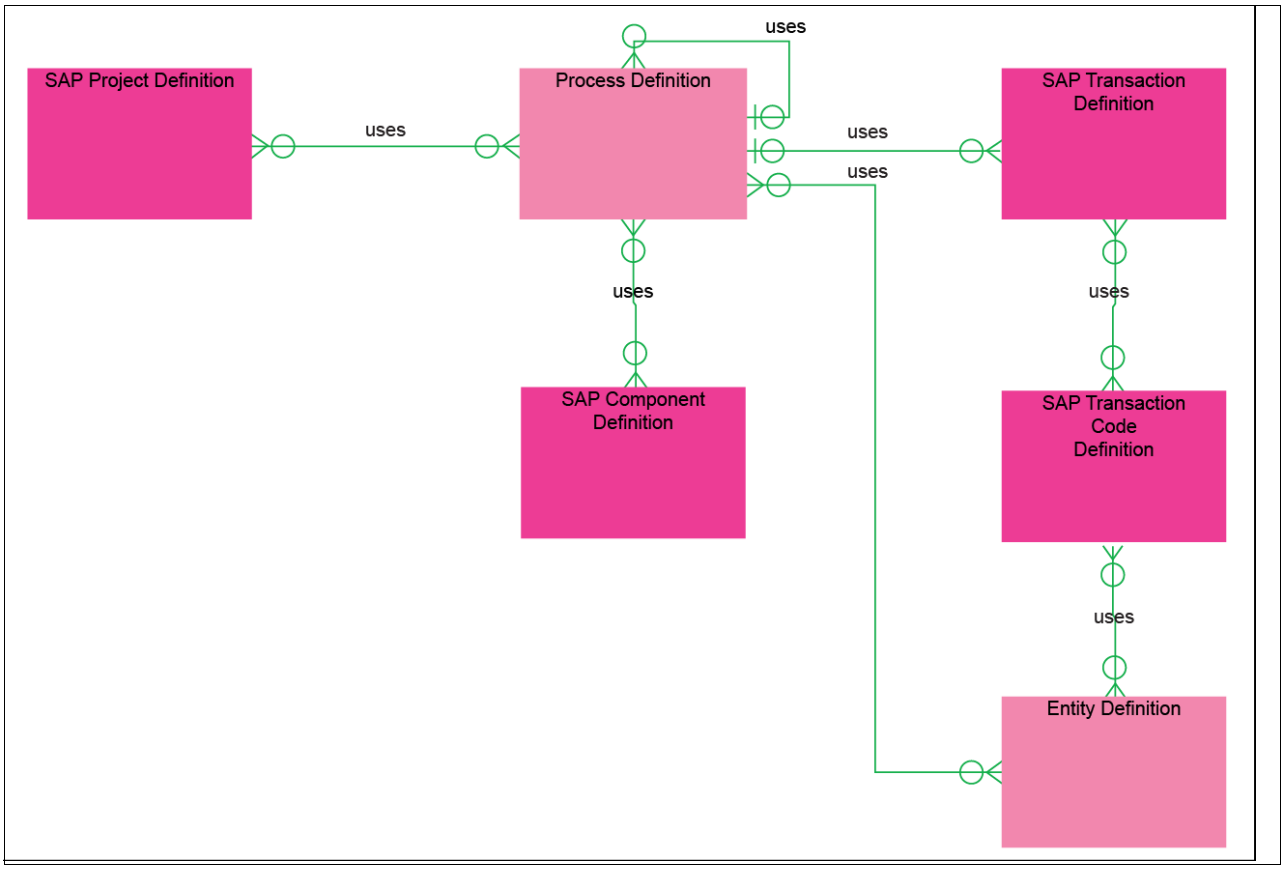

Figure 10-17 Sample model of the SAP Solution Manager component architecture

With continuous business planning, organizations can generate a single, holistic view of the high-level processes that, within SAP Solution Manager, are defined in detailed logic flows that incorporate manual process *and* processes supported by other systems. In addition, the comprehensive integrated environment and analysis tools of the solution help to quickly identify and analyze any problems that might occur.

By effectively incorporating the SAP landscape into an enterprise architecture, organizations gain better insight into their overall technology plan. Continuous business planning for SAP enables SAP customers to plan, manage, and control change to their SAP environment. Businesses become ready to respond to change.

Visualizing an integrated view of SAP projects, blueprints, and landscapes within the context of the enterprise architecture, business processes, data, organization, and roles within a business-process context enables companies to achieve the following capabilities:

- Compare *as-is* and proposed *to-be* solutions.
- Automate synchronization of models between SAP Solution Manager and Rational System Architect. This approach minimizes manual entry of business architecture information from SAP Solution Manager.
- Access individual SAP process objects as required when building the business process models managed by Rational System Architect. This approach delivers the benefits of comprehensive business enterprise architecture to SAP environments.

- Identify interdependencies and execute impact analysis.
- Record a history of business process models and their changes. Validate that required changes were made correctly, and that no unauthorized changes were applied, by tracking changes made to business process models.
- Compare new and old models to help identify and validate revisions. This approach eliminates manual copying errors when replicating models.
- Uniquely identify objects so that the applications can continue to recognize objects even if their names are changed. This approach ensures that the designed architecture is the architecture implemented.

10.7 Summary

To help SAP customers reduce cost and risk, IBM and SAP have come together to help customers overcome the challenge of managing risk and change to SAP landscapes. The solutions that make up the IBM DevOps for SAP approach are designed to achieve the following objectives:

- Application lifecycle management for SAP
- Collaborative development for SAP
- Continuous testing for SAP
- Continuous release and deployment for SAP
- Continuous monitoring for SAP
- Continuous business planning for SAP

Figure 10-18 shows the IBM Rational tools that implement the five core Rational solutions for continuous SAP delivery.

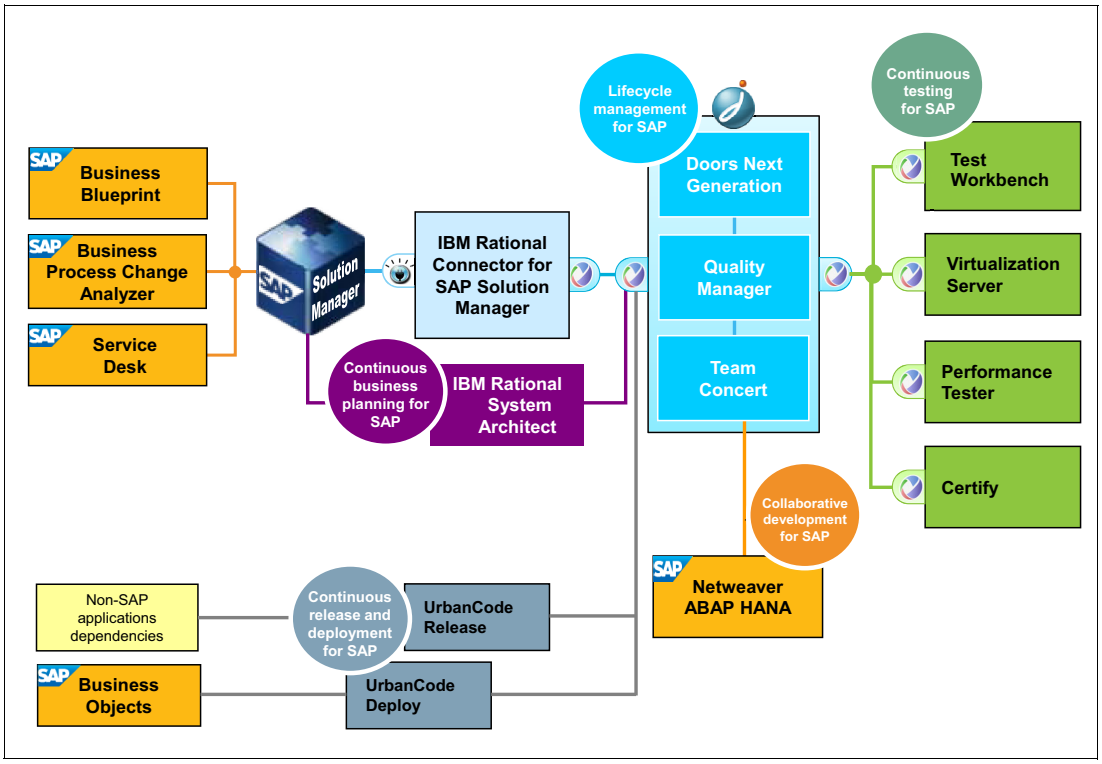

Figure 10-18 IBM DevOps for SAP: Tools architecture

IBM DevOps for SAP enables all participants (business teams; architects, developers and testers; outsourced partners; and IT operations and production) to engage throughout the SAP implementation and change lifecycle, and to align themselves with a common goal fueled by continuous delivery and shaped by continuous feedback.

Thanks to the close collaboration between IBM and SAP, SAP customers benefit from extensive integration and a powerful blend of IBM and SAP leading practices and technology that help to optimize and connect the following aspect of development for both SAP and non-SAP applications:

- Enterprise planning
- Application lifecycle management
- Quality management
- Development
- Testing
- Change management
- DeploymentS

10.8 References

These websites are also relevant as further information sources:

- IBM Rational solutions for SAP

  http://www.ibm.com/software/rational/solutions/sap/

- Managing change in SAP: reducing cost and risk with IBM DevOps

  http://www.ibm.com/common/ssi/cgi-bin/ssialias?subtype=FY&infotype=PM&appname=SWGE_RA_VF_WWEN&htmlfid=RAF14154WWEN

11

Systems security for SAP

With respect to information technology (IT) security, an essential task is for enterprises to carefully plan and manage SAP environments in the context of their end-to-end architecture. SAP systems offer many security capabilities. However, some enterprises find that they need enhanced security in the SAP system, and many have requirements to integrate SAP security with other, non-SAP systems in the enterprise. IBM can help clients meet these needs.

IBM provides proven, integrated, scalable, and end-to-end enterprise security products. The offerings span the full spectrum of the IBM Security Framework, composed of governance, risk management, and compliance (GRC), security intelligence, people, data, applications (apps), and infrastructure.

This chapter provides information about the objectives and components of the architecture defined to support security management technology integration as part of a larger reference architecture for SAP implementations. This chapter describes technologies and products available from IBM to integrate with SAP systems and applications to provide the inclusive, end-to-end security that enterprises require.

This chapter includes the following topics:

- 11.1, “SAP systems and IBM security management integration overview” on page 276
- 11.2, “SAP systems security and IBM security management integration scenarios” on page 279
- 11.3, “Identity system scenarios” on page 281
- 11.4, “Authentication system scenarios” on page 285
- 11.5, “Authorization system scenario” on page 291
- 11.6, “Audit system scenarios” on page 293
- 11.7, “Identity management products and solutions” on page 297
- 11.8, “Access management products and solutions” on page 299
- 11.9, “Audit products and solutions” on page 302
- 11.10, “References” on page 305

© Copyright IBM Corp. 2014, 2015. All rights reserved.

11.1 SAP systems and IBM security management integration overview

When SAP solutions constitute the backbone of the enterprise, application security, risk, and threat management must encompass SAP systems and applications. The consequences of a breach are extremely negative given the business-critical nature of SAP applications, such as financial operations, human resources, supply chain, or customer relationship management (CRM).

SAP applications run at the core of an enterprise and need particular protection. SAP solutions are also an integrated system and part of the IT infrastructure. Therefore, any errors might have a widespread effect. Security of enterprise resource planning (ERP) systems must provide protection at several layers:

- User layer
- Communications layer
- Application layer
- Database and data storage layer

The main goal is to integrate SAP solutions into an enterprise security architecture to accomplish a comprehensive view and control mechanism for the entire heterogeneous IT landscape.

A security management solution that integrates with SAP systems and applications must address the following non-functional requirements:

- Access to critical and sensitive activities in SAP applications is controlled.
- User access is based on user role and responsibilities.
- Authorization within roles is based on the role definition.
- User provisioning and role management policies and procedures are implemented.
- Logon procedures are reduced.
- Effective controls and monitoring are in place.

Figure 11-1 on page 277 illustrates the reference architecture and generic components for integration of SAP systems and applications into an enterprise environment. Common security components can be provided by corporate infrastructure or by integration deployment.

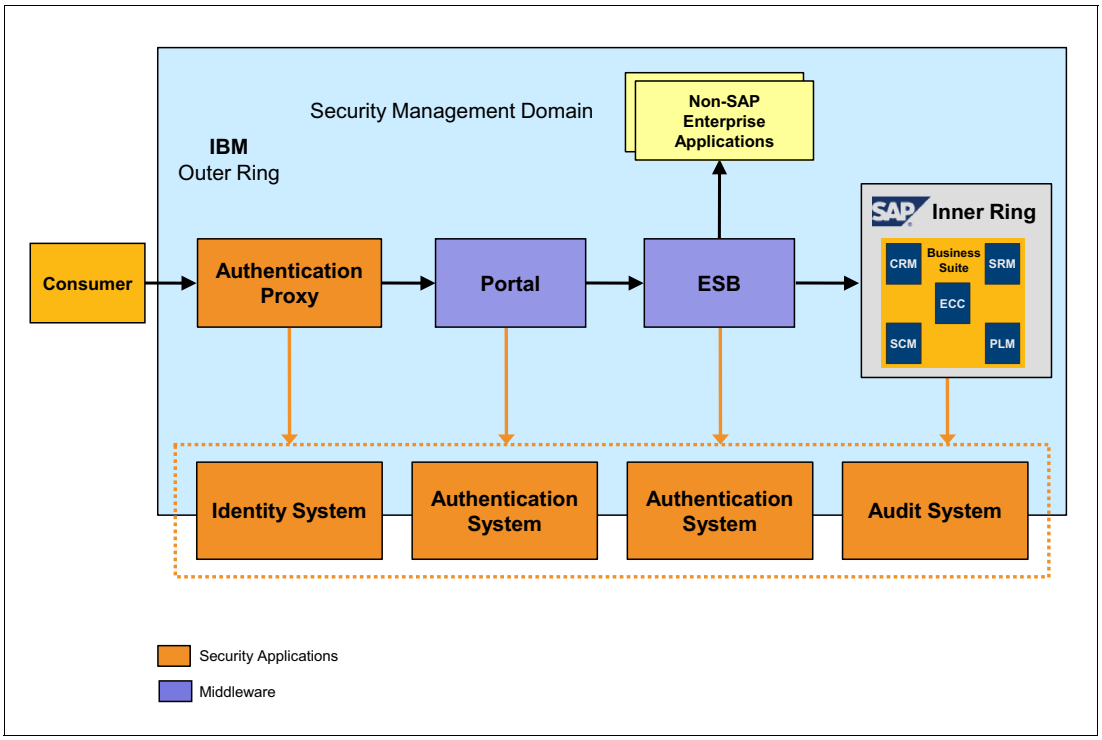

Figure 11-1 Generic security reference architecture

11.1.1 Key capabilities: Logical components of security reference architecture

Table 11-1 on page 278 lists the security components that are depicted in Figure 11-1 and their associated functionalities. The positioning of the fundamental components follows the general IBM Reference Architecture for SAP software principles with an inner ring (core SAP systems and applications), and an outer ring that consists of best-in-class IBM security solutions. Specific SAP software integration capabilities ensure that SAP-specific requirements are addressed.

Security management is an enterprise-wide task and should be organized accordingly. For example, organizations should establish company-wide responsibilities, such as a Chief Security Officer, and introduce and maintain general security policies and procedures.

Supporting technology should meet the enterprise approach, with security management systems that share functionality across enterprise applications (outer ring and inner ring), and that are used to enforce rules and regulations, constantly control conditions, and intervene in case of deviations or incidents.

The following major security management components are shown in Figure 11-1:

- Identity system
- Authentication system
- Authorization system
- Audit system

Table 11-1 shows the general functionalities provided by those systems, sample scenarios, and use cases.

Table 11-1 Security management components and general functionalities and use cases

Security management system	General functionality	Sample scenarios and use cases
Identity system	► Manage identities centrally in an efficient manner. ► Provision identities consistently to the respective components.	► Identity feed ► User provisioning ► User account recertification ► Password management ► User self-service
Authentication system	► Prove the identity of clients that issue requests to back-end components. ► Map identities across different components. ► Provide multiple enforcement points to use unified decision points.	► Access management ► Single sign-on
Authorization system	► Apply business policies consistently across multiple policy enforcement points to provide fine granular access to business functions. ► Enable standard-based unified policy management across different vendor components.	► Policy-based identity and access governance. ► Systems can be used only by appropriate users. ► Usage of appropriate functions or system resources is enabled only for authenticated users.
Audit system	► Establish a heterogeneous auditing capability to control all relevant enterprise components to ensure consistency and compliance.	► Security monitoring (user login logging, illegal act monitoring) ► Data protection and redaction ► Source code analysis and application security testing

11.1.2 Mapping logical security architecture components to IBM and SAP software

In Table 11-2, the security logical architecture components introduced in Figure 11-1 on page 277 are listed and mapped to the related IBM and SAP software tools that provide the required capabilities.

Table 11-2 Generic components and product mapping

Generic components	Software
Consumer	► Service requester ► Rich client ► Fat client, such as the SAP graphical user interface (GUI) ► Internet browser
Proxy	► IBM Security Access Manager for Web ► IBM WebSphere proxy ► IBM WebSphere DataPower Family
Portal	► IBM WebSphere Portal Family ► SAP Enterprise Portal

Generic components	Software
Middleware Enterprise service bus (ESB)	► IBM WebSphere Message Broker ► IBM WebSphere DataPower
Back-end (SAP)	► SAP Business Suite ► SAP ERP (ERP Central Component, ECC) ► SAP CRM
Identity system	► IBM Security Identity Manager ► SAP Identity Management (Central User Administration (CUA), and so on) ► IBM Security Directory Server
Authentication system	► IBM Security Access Manager for Web ► IBM Security Access Manager for Mobile ► IBM Tivoli Federated Identity Manager
Authorization system	► IBM Tivoli Security Policy Manager
Audit system	► IBM QRadar® Security Intelligence Platform ► IBM InfoSphere Guardium® ► SAP BusinessObjects Access Control (SAP GRC) ► Virtual Forge CodeProfiler for IBM Security AppScan®

11.2 SAP systems security and IBM security management integration scenarios

This section has more detail about security management components and common scenarios.

11.2.1 Key solution components

To better understand the sample scenarios, first describing the key solution components is important.

An integrated SAP system landscape is a complex SOA-based environment, which consists of various components that interact with each other and with external components. A generic security framework helps to achieve a unified security view of packaged business applications, middleware, and user interface (UI) components.

Solution components must satisfy a variety of security requirements:

- Manage identities across various components.
- Define access and authorization policies in a heterogeneous environment.
- Provide common functions, such as single sign-on (SSO), within the system landscape.
- Provide auditing capabilities.
- Ensure compliance with corporate policies.

11.2.2 Generic components and concepts of a security architecture

A security architecture must include decision points (DP) and enforcement points (EP) to evaluate and control access to resources across the entire IT infrastructure (both SAP-based and non-SAP-based infrastructures).

As defined in International Organization for Standardization (ISO) 10181-3, the following list describes generic components participating in an access request process (although using different terminology):

1. The initiator requests access to a target resource.
2. The request is intercepted by the policy enforcement point (PEP).
3. The PEP submits the request, with information about the user, to the policy decision point (PDP).
4. The PDP returns a decision (Boolean response) to the PEP.
5. The PEP enforces the decision.

Figure 11-2 shows the access request flow as defined in ISO 10181-3.

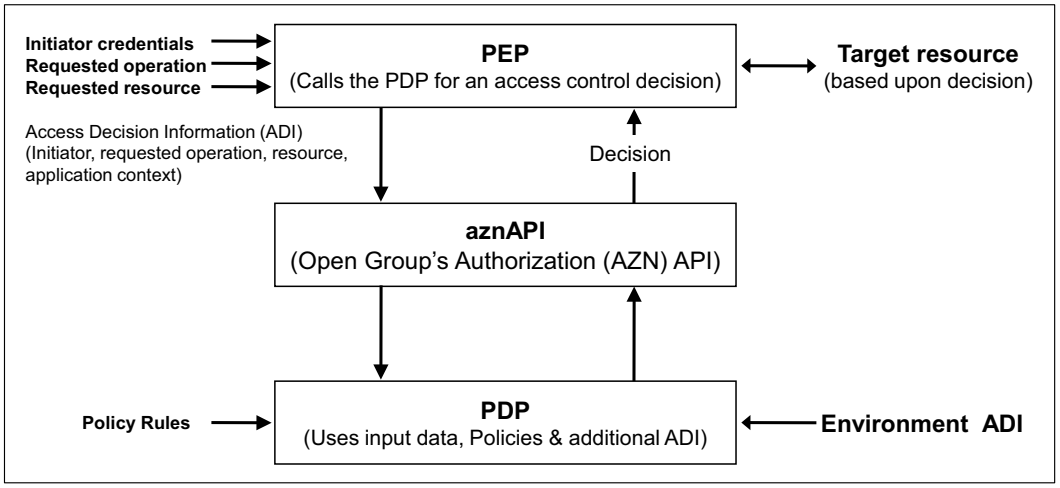

Figure 11-2 Components participating in an access request as defined in ISO 10181-3

When designing a security architecture, consider these other concepts and components:

- Introduce several security-relevant areas:
  - Entry zone, enforcement, and DMZ
  - Intermediate zone, middleware, and integration
  - Back-end zone and business applications
  - Security management zone
- Focus on using existing products rather than building custom code into integration artifacts: Avoid custom mapping and transformation.
- Use existing corporate security components: Security-as-a-service.
- Consistency across various components:
  - Same identity, authentication, policy, and audit systems.
  - One management layer that serves various components.

- Achieve interoperability by supporting open standards:
  - Security Assertion Markup Language (SAML)
  - X.509, Web Services Security (WS-Security)
  - Extensible Access Control Markup Language (XACML)
  - WS-Trust
  - WS-Policy
  - WS-SecurityPolicy
  - WS-MetaDataExchange
  - WS-Transfer

In addition, SAP-specific security integration challenges exist:

- Embed packaged application (SAP) into an SOA environment:
  - SAP business application offers integrated authentication, authorization, and policy schema.
  - SAP business application supports only limited externalization of these core features.
- SAP-based business services can be exposed across the enterprise:
  - Achieve a generic enforcement strategy across domains.
  - Accountability and audit to track users activities.

11.3 Identity system scenarios

The identity system provides the following main functionalities:

- Manage identities centrally in an efficient manner.
- Provision identities consistently to the respective components.

The following sections describe identity system sample scenarios. For information about Identity management products available from IBM, see 11.7, “Identity management products and solutions” on page 297.

11.3.1 Identity management scenario

IT security requires processes, methods, and tools to manage the identities that are used to control access to resources.

Identity management encompasses all of the data and processes that are related to the representation of an individual involved in electronic transactions.[1] It is the process that enables business initiatives by efficiently managing the user lifecycle. including identity and resource provisioning for people (users), and by integrating it into the required business processes.

For user administration and enterprise identity lifecycle management, IBM provides IBM Security Identity Manager (formerly known as Tivoli Identity Manager). It enables corporate identity management, and integrates SAP systems and applications for user provisioning. This product enables you to separate identity groups, such as intranet users, or external identities (business partners, suppliers, and so on).

[1] *Identity Management Design Guide with IBM Tivoli Identity Manager*, SG24-6996.

Figure 11-3 shows the interactions of identity management components.

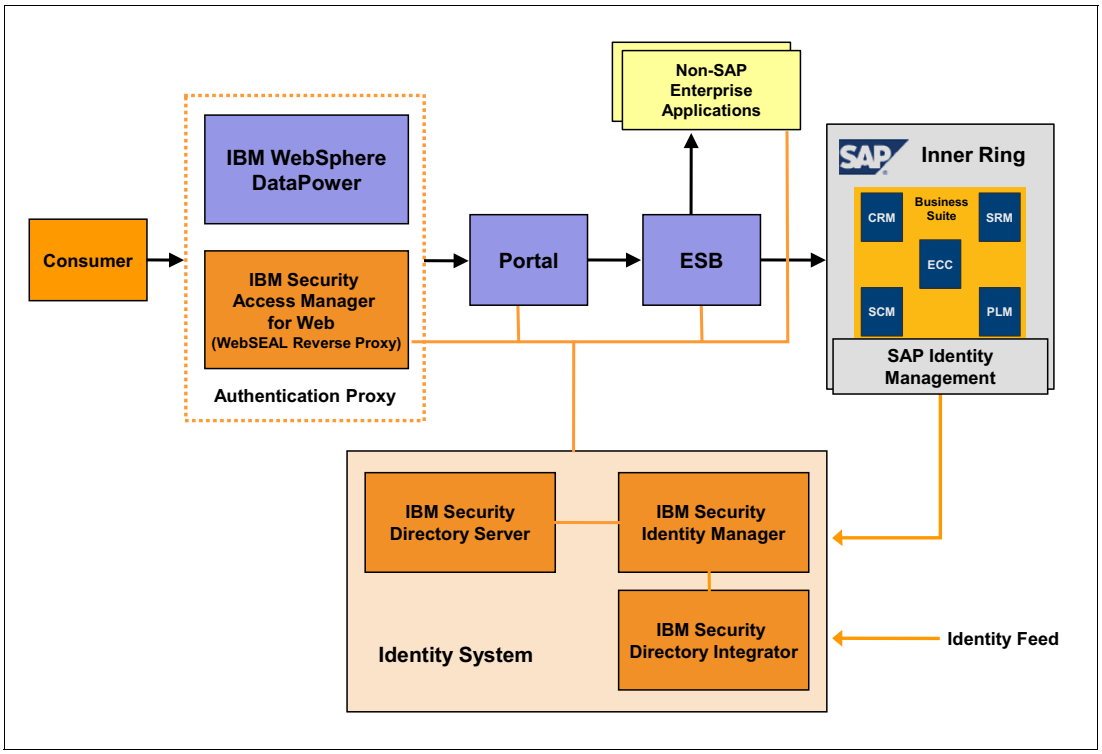

Figure 11-3 Identity management components interaction

The key identity management components are shown in the box labeled Identity System in Figure 11-3. These components are the foundation layer upon which runtime authentication and authorization is provided by the IBM Security Access Manager for Web reverse proxy. The following list describes the Identity System:

- IBM Security Directory Server. Forms the identity data store. Organization identities, accounts, and associated entitlements, such as roles and groups, are stored here. This component is referenced by the authentication proxy at run time.
- IBM Security Identity Manager. Provides the central location from which users and administrators can request and provision access to business applications, such as SAP applications. Supports request and role-based access workflows for account provision. It also provides account auditing and recertification functions to help ensure that minimal rights are maintained for each user. It uses a directory, such as IBM Security Directory Server, to persist account and identity data.
- IBM Security Directory Integrator. This component provides the business application (for example, SAP applications) connectivity capability to IBM Security Identity Manager. IBM Security Directory Integrator is used to ensure that accounts and account data in SAP software are synchronized and up-to-date with data managed in IBM Security Identity Manager. This process is known as *account reconciliation*.

  IBM Security Directory Integrator is used to push or provision accounts from IBM Security Identity Manager as a result of the provisioning workflow. IBM Security Directory Integrator is also used to pull human resources (HR) data from SAP applications into IBM Security Identity Manager. This process is known as an *identity feed*.

With the introduction of SAP NetWeaver Identity Management, SAP closed a gap in their product portfolio by providing a tool that enables maintaining user information that spans SAP systems, including both Advanced Business Application Programming (ABAP) and Java platforms. SAP NetWeaver Identity Management is offered as an additional installation to ABAP and User Management Engine (UME) local user administration and CUA. It also provides functionality to integrate third-party identity management tools and non-SAP apps.

SAP offers SAP NetWeaver Identity Management services and interfaces for partners to implement solutions, enabling the integration of heterogeneous environments.SAP NetWeaver Identity Management includes a core set of integration adapters for Lightweight Directory Access Protocol (LDAP) and relational database stores. However, it relies on third-party partners for integration adapters to other data stores. For more information, see 11.7, “Identity management products and solutions” on page 297.

11.3.2 Identity feed scenario

An *identity* is the subset of profile data that uniquely represents a person in one or more repositories, and additional information related to that person. For example, an identity might be represented by a unique combination of a person's first, last, and full name, and employee number. Additional information might include phone numbers, manager’s name, and email address. A data source can be different, depending on the business scenario:

- For Internet scenarios, the data source can be a self-service user interface.
- For federated scenarios, user data is often provided by the business partner.
- For corporate usage, the data source can be a user repository or a file, directory, or custom source.

User data in a corporate environment is often provided by the HR system. In the case of SAP, this is the SAP ERP Human Capital Management application (SAP HCM, formerly known as SAP HR).

Reconciliation for an identity feed is the process of synchronizing the data between the data source and the identity management system. The initial reconciliation populates the identity management system with new users, including their profile data. A subsequent reconciliation creates new users and also updates the user profile of any duplicate users that are found.

The decision about what option to use for the identity feed depends on available systems for authoritative identity sources, and the level of trust associated with them. Table 11-3 lists the identity feed architectural decision options and selection criteria.

Table 11-3 Identity feed architectural decision

Subject	Identity feed
Architecture decision	Authoritative identity sources and identity master, and process to synchronize identity data
Alternatives	► HR-driven identity feed, such as SAP HCM ► User self-service registration, such as web identity ► Enterprise directory, such as LDAP ► Application master, such as SAP ERP ► Combination of previously listed alternatives
Decision criteria	► Business scenario requirements, such as business-to-business (B2B), business-to-consumer (B2C), and so on ► Availability and trustability of data sources ► Accessibility and format of the data ► Initial feed or periodic synchronization

Subject	Identity feed
Decision examples	► When employee information is maintained in HR, availability and level of trust are high. As an illustrative example, it is possible to assemble solutions where SAP HCM data can be extracted and updated automatically using the SAP Business Object Repository interface. ► Specialized systems can be used as authoritative sources of specific data, for example, use IBM Notes for email addresses, telephone directory for phone numbers, and so on.
Architecture	Use a synchronization component for the periodic identity feed into the identity management system, for example, access to SAP HR (person data) with IBM Security Directory Integrator, and merge with additional information from other systems (for instance, email and phone number).

11.3.3 User provisioning scenario

SAP-specific user administration tasks can be performed by local systems using specific applications, such as SAP UME, which is primarily the user store for the SAP Enterprise Portal. SAP CUA can be set up as the user administration hub for SAP NetWeaver ABAP-based applications.

SAP software provides integration connectivity tools and application programming interfaces (APIs) to enable third-party tools to access user repositories and management functions. The IBM security products make use of these interfaces, and partly extend their functionality to integrate SAP systems and applications into a corporate-wide identity management system. Figure 11-4 shows SAP user provisioning options for the various SAP platforms.

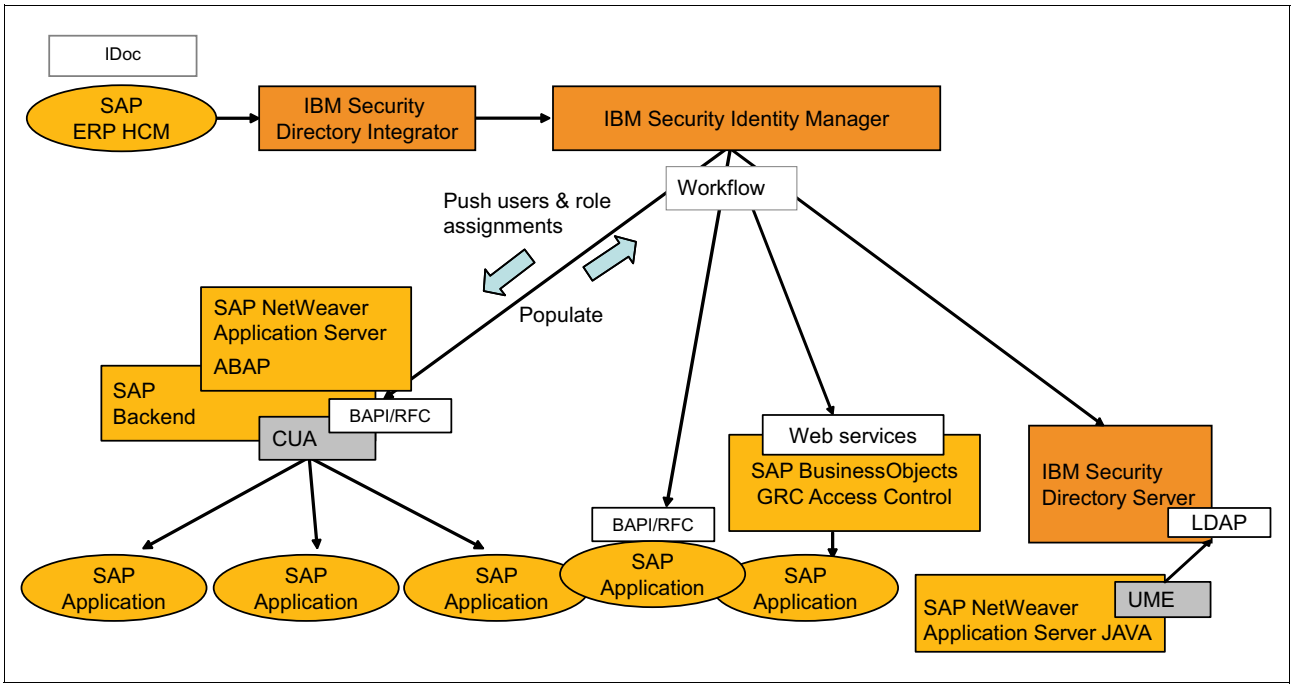

Figure 11-4 SAP user provisioning options

The decision about which option to use, and how it fits into the overall identity management system, depends on several factors, such as the SAP applications and platforms that are used. Also consider whether separation, pooling, or additional compliance management requirements exist. Table 11-4 lists the user provisioning architectural decision options and selection criteria.

Table 11-4 User provisioning architectural decision

Subject	User provisioning
Architecture decision	Integrate and pool SAP systems as provisioning targets for IBM Security Identity Manager.
Alternatives	► Direct user provisioning through SAP ABAP clients ► User provisioning to SAP clients using SAP CUA ► User provisioning to SAP clients using SAP GRC solutions ► User provisioning to SAP Java systems using LDAP
Decision criteria	► SAP system separation requirements, for example development, quality assurance (QA), production, training, and so on ► Enterprise directory usage ► Usage of SAP CUA required ► Use of SAP GRC solutions for separation of duties (SOD) verification
Decision example	► When the corporate LDAP system is used for the SAP Enterprise Portal persistence store, the user provisioning should go to that system for the SAP application users. ► When detailed SOD verification is required for SAP systems, it should be integrated into the provisioning workflow.
Architecture	Use IBM Security Identity Manager as the hub for provisioning users to SAP environment targets.

11.4 Authentication system scenarios

The authentication system provides the following main functionalities:

- Prove the identity of clients that issue requests toward back-end components.
- Map identities across different components.
- Enable multiple EPs to use unified DPs.

The following sections describe authentication system sample scenarios. For information about access management products available from IBM, see 11.8, “Access management products and solutions” on page 299.

11.4.1 Access management scenario

IBM provides IBM Security Access Manager for Web (formerly known as IBM Tivoli Access Manager) and other products in the IBM Security Access Manager Family to secure enterprise environments. They might be deployed to secure internal IT infrastructures and extranet resources. The IBM Security Access Manager products provide user authentication and authorization. For federated SSO solutions, advanced identity mapping capabilities are supported. For example, intranet ID and Internet ID mapping.

IBM Security Access Manager for Web acts in conjunction with WebSphere DataPower as the web point of contact and EP.

One of the benefits of using a central *access management* component to secure back-ends and heterogeneous environments is that it externalizes authentication and, therefore, the application components are not required to perform authentication processes and mechanisms. By supporting several authentication protocols and login ticket technologies, the IBM access management products support a variety of systems and applications:

- WebSphere-based applications using the Lightweight Third Party Authentication (LTPA) generator
- Middleware components using the Security Token Service (STS)
- Identification ID transformation of intranet or extranet IDs into SAP ID

The authentication system provides the following key functionalities:

- Secure user interaction (browser) and business application interaction (web services) across components
- End-to-end lifecycle management enabled by establishing service provider (SP) and Identity Provider (IdP) patterns

Figure 11-5 shows the interaction of the access management components.

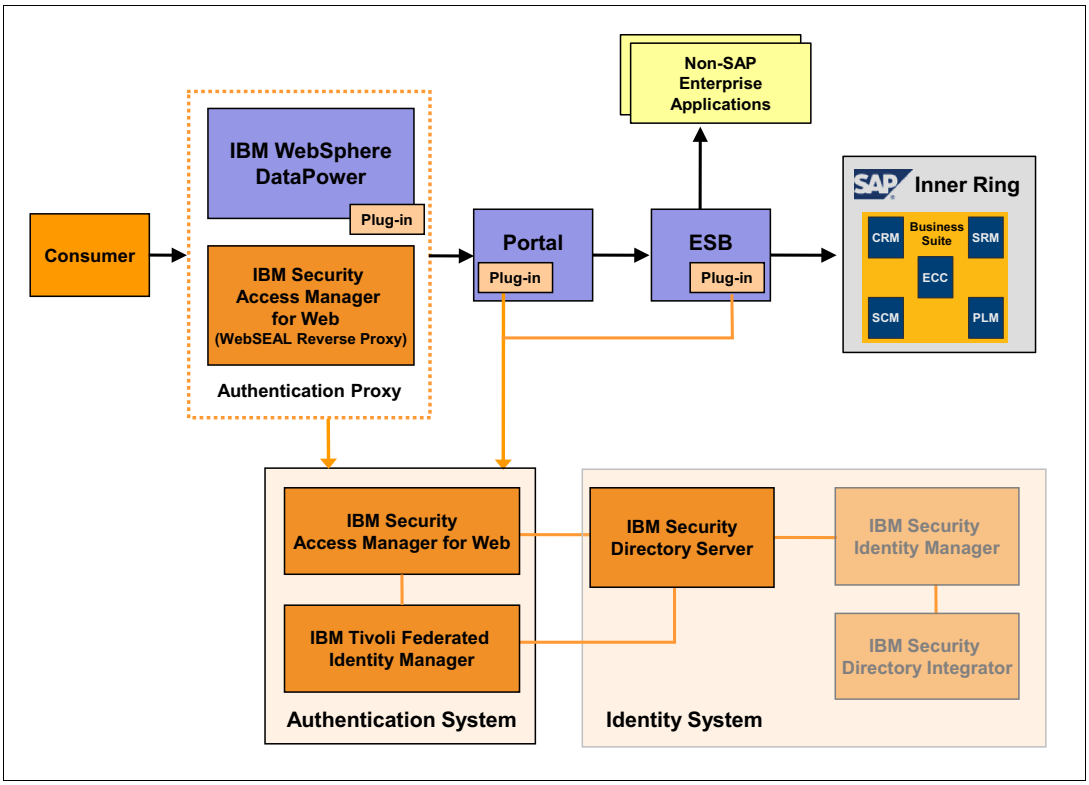

Figure 11-5 Access Management components

The following enforcement point components are used for the security architecture for SAP:

- WebSphere DataPower is the security and PEP for all web services-based traffic into the SAP system.
- The IBM Security Access Manager for Web reverse proxy is the security and PEP for all HTTP-based traffic (UI consumption) into SAP and non-SAP systems and applications.

Establishing these two control points enables the usage of common security services from identity management, access management, policy management, and common audit mechanisms:

- IBM Security Access Manager for Web is the centralized authorization DP that provides a consistent authorization policy across the corporation.
- WebSphere DataPower is a PEP.
- A data-centric authorization scheme, using IBM Security Access Manager for Web object space and authorization policies, provides interface-independent solutions (they function whether the interface is based on web services SOAP, REST interface, and so on).
- Object space based on the Enterprise Subject Area Model (ESAM) definition, which is a high-level cross-process information model representing all of the categories of information within the enterprise. It is the foundation of the interface-independent authorization solution.
- WebSphere DataPower authentication, resource mapping, authorization, and audit architecture provide security at the boundary of the infrastructure, and at the front end of applications.
- The web service interface with Organization for the Advancement of Structured Information Standards (OASIS) XACML standard-based protocol between IBM WebSphere DataPower and IBM Security Access Manager provides the full IBM Security Access Manager for Web PDP capabilities to WebSphere DataPower (PEP).
- IBM Security Access Manager for Web authorization rules and server plug-ins implement complex authorization requirements.

For more information, see 11.8, “Access management products and solutions” on page 299.

11.4.2 Single sign-on (SSO) scenario

Single sign-on is part of enterprise security, and it is an asset outside the scope of the SAP transformation program. However, the SAP transformation program must consider enterprise SSO and integrate with it.

Generally, an SSO architecture reduces the number of logins, ideally to just one, either by authenticating a verified user from one application to another or by authenticating on behalf of the application by an authentication proxy. In addition, IBM distinguishes between desktop applications, web environments, and federated systems to approach access management and SSO. Table 11-5 list the SSO architectural decisions alternatives and selection criteria.

Table 11-5 Single sign-on architectural decision

Subject	Single sign-on
Architectural decision	Implement an authentication system that meets certain computing model and technology criteria.
Alternatives	► Desktop or enterprise SSO. A mechanism to store and manage a user’s account IDs and their associated password or other credentials. ► Web SOS. A mechanism by which HTTP services are aggregated behind the security proxy. The security proxy is responsible for establishing authenticated sessions in a centralized point in the infrastructure. ► Federated SSO. A mechanism by which users can be asserted from one administrative domain to another.

Subject	Single sign-on
Decision criteria	► Authenticate user against an application-specific repository. The user population is contained within the enterprise. Use of desktop applications, proprietary communication and protocols, corporate (private) networks (for example, use of SAP GUI, SAP Supply Network Collaboration (SNC), or SAP Login Ticket). ► Authenticate a user against an application-specific or enterprise repository. User population expands to the Internet. Use of standards-based web protocols (HTTP Secure), private or public networks, B2B, and B2C (for example, SAP Enterprise Portal). ► Authenticate a user against an enterprise repository. Also authenticate users and other network entities (for example, web services) originating from foreign companies. The user population expands to the Internet and other enterprises. Use of public networks, standards-based protocols (web services, SAML), B2B (for example SAP Enterprise Portal), SAP Enterprise Services (SOA) (for example, with SAP Process Integration, SAP Application Server is the WS Provider, and so on).
Decision example	► When the SAP environment is accessed mainly with the SAP GUI, use IBM Security Access Manager for Enterprise SSO. ► When the SAP environment is accessed through DMZ and a browser (user authentication), use IBM Security Access Manager for Web. ► When the SAP environment is accessed by Internet and different channels that require token mediations, use IBM Tivoli Federated Identity Manager.
Architecture	Hub and spoke architecture with a central authentication component configured to be trusted by SAP back-ends. SSO technologies can be combined in a mixed environment, such as deployment of an authentication proxy with identity federation.

Figure 11-6 on page 289 illustrates an example of how SSO based on IBM Security Access Manager for Web reverse proxy can be integrated with SAP in a web environment.

The integration is based on user authentication using reverse proxy, and obtaining a security identity token in a format that is consumable by SAP or SAP partner applications. Because SAML is the mechanism preferred by SAP for third-party logon, this scenario is designed using a federated identity management approach that enables the use of SAML for authentication and the use of the Security Token Service (STS) provided by Tivoli Federated Identity Manager.

If different user IDs are used, the primary authentication is to the reverse proxy as the point of contact. As part of the authentication process, Tivoli Federated Identity Manager is able to perform the mapping of the authenticated user name to the SAP user name for the authenticated user. This mapped user name is then forwarded to the SAP environment. Tivoli Federated Identity Manager mapping methods are flexible. A common method includes a lookup of the required attributes from the LDAP directory.

Restriction: Tivoli Federated Identity Manager supports SAML, however, IBM does not provide a formally supported integration based on SAML for SAP environments. The integration is technically feasible, but IBM is not committed to accept support authorized program analysis reports (APARs) if problems occur.

For more information, see the article *Integrating IBM Federated Identity Manager 6.2.2 with SAP Login Tickets*, which shows an example of an unsupported integration. The article describes how the Token Security Service in Tivoli Federated Identity Manager V6.2.2 can be integrated with the SAP Login Ticket to validate user identity. The article is available at the following location:

`http://ibm.co/1l2xlpV`

Figure 11-6 shows SSO integration with SAP. See 11.4.3, “Identity propagation scenario” on page 289 for details.

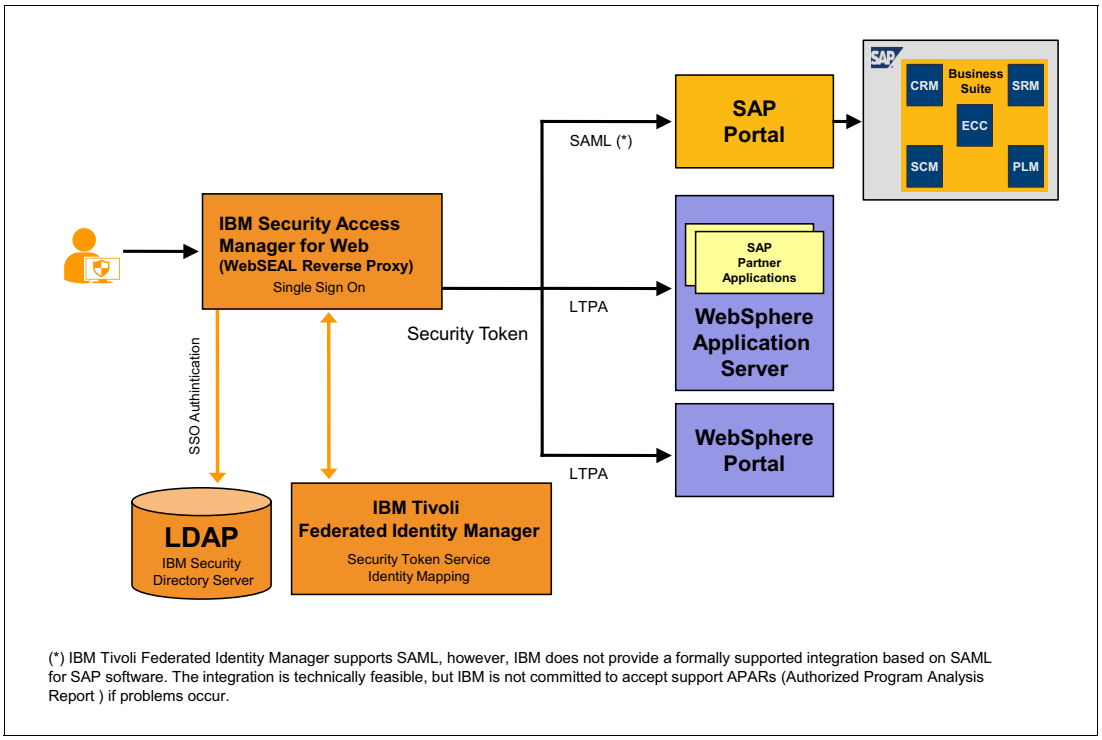

Figure 11-6 Web-based SSO to SAP applications

11.4.3 Identity propagation scenario

Besides policy-based user provisioning at the enterprise level, with a flexible workflow and a broad range of supported systems to manage, the existence of a variety of user identities requires mapping and propagation at run time. Where possible, 1:1 mapping of user name to SAP user name is the simplest approach. Alternatively, another approach is to use the IBM Tivoli Federated Identity Manager alias service to store the mapping from the local user name to the SAP user name.

A Tivoli Federated Identity Manager mapping rule can be used to query the pre-provisioned SAP user name from the IBM Tivoli Federated Identity Manager alias service. The SAP user name is typically provisioned separately by an identity management product, such as IBM Security Identity Manager.

The Tivoli Federated Identity Manager alias service has helper APIs that can be called from a mapping rule. The Tivoli Federated Identity Manager alias service can be configured to be either in LDAP, or directly in a relational database using Java Database Connectivity (JDBC). Both models are supported with Tivoli Federated Identity Manager.

The following scenarios require identity mapping and propagation:

- Different user identities passing between components.
- Identity has to be mediated (Tivoli Federated Identity Manager provides STS).
- Back-end applications, for example SAP applications, need user identity.

In these scenarios, the identity master provides user ID information to be mapped or propagated to the target systems and applications, such as identity synchronization (see Table 11-6).

Table 11-6 Identity propagation architectural decision

Subject	Identity mapping and propagation
Architecture decision	Use of components and flow of identities for specific use cases.
Alternatives	► The ESB performs identity mediation. ► Authentication proxy calls STS.
Decision criteria	► Message flow is web services only ► User-based authentication ► Options of available trust configurations, for example, when WS provider can handle one trust relationship only ► Complexity of user ID mapping, for example, 1:1 mapping, mapping table, or composition formula ► Use of authentication formats and general token variants for different components in the message flow, for example, user certificates, user name token, LTPA, SAML, and so on ► Use of SAML token variants, such as Holder of Key (requires token mediator) or sender-vouches
Decision example	► Web Services scenario using WebSphere Message Broker and WebSphere DataPower between WS consumer and WS provider (SAP application server) and Tivoli Federated Identity Manager (as the Token Mediator.) Trust relationship between Security Token Service (STS) (IBM Tivoli Federated Identity Manager) and WS provider (SAP). ► WebSphere DataPower is both, ESB and token mediator to issue and send SAML tokens to SAP back-ends. Trust relationship between WebSphere DataPower and SAP software. **Note:** These examples are for illustrative purposes only. These integrations are not formally supported by IBM.
Architecture	Use specialized components to handle multiple variants of authentication mechanisms (authentication proxy), message flows (ESB), token formats, and trust relationships (STS).

WebSphere DataPower can prove identity using a corporate user directory (corporate LDAP as provided by IBM Security Directory Server), and request an identity mediation from Tivoli Federated Identity Manager STS when necessary to call the intermediate ESB service. WebSphere Message Broker can trust the calling identity or request again an identity mediation from Tivoli Federated Identity Manager STS to call the SAP back end using the application and UserID.

Figure 11-7 shows the identity propagation components and flow.

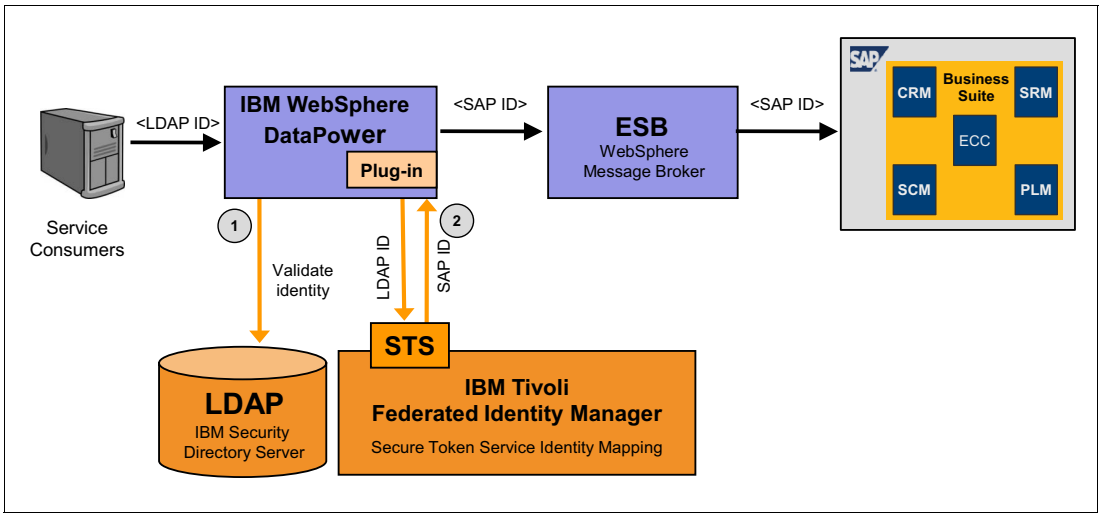

Figure 11-7 Sample identity propagation components and flow

11.5 Authorization system scenario

Tivoli Security Policy Manager can be used as a central policy management component. Tivoli Security Policy Manager enables applications to externalize their security functions, to centralize a runtime engine driven by a standards-based policy framework.

This policy can be centrally managed, is external to the application, and can be updated independent of application development cycles. With this setup, management of the security policy and associated security infrastructure can be removed from the application logic.

Other security components, such as IBM Security Access Manager for Web, WebSphere DataPower, or WebSphere Message Broker components, enable the usage of external security policy management frameworks, such as Tivoli Security Policy Manager (see Figure 11-8 on page 292).

The authorization policy for access and entitlements should include the following requirements:

- Make early access decision based on common policy.
- Centralize DP for access and entitlements.
- Avoid inconsistency in managing policies across different enforcement points.

Figure 11-8 shows the interactions between authorization system components.

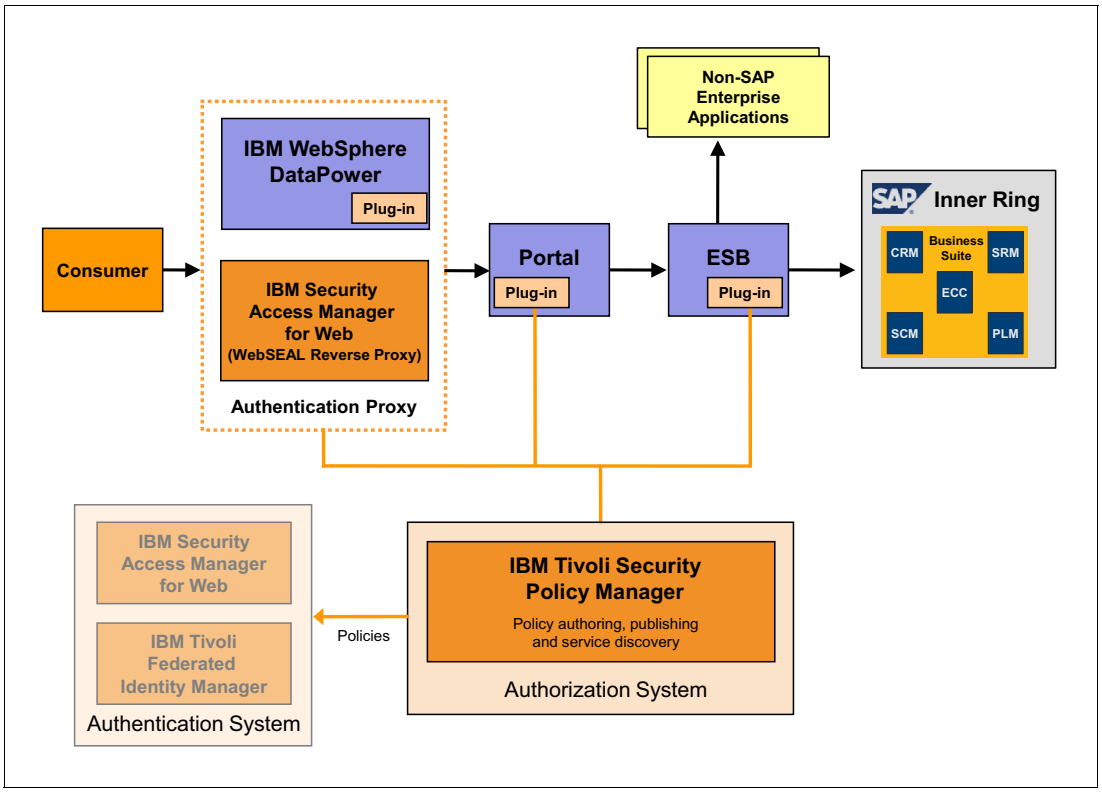

Figure 11-8 Authorization system components interaction

Tivoli Security Policy Manager centralizes security policy management and fine-grained data access control for applications, databases, portals, and services.

Tivoli Security Policy Manager uses the XACML standard and provides an implementation of the processing model that the standard defines. Specifically, it implements the policy administration point (PAP), PEP, and PDP components for a variety of systems (such as, WebSphere DataPower Appliances), databases, and application servers. Although the SAP software does not support XACML or similar standards to externalize security policy management, an alternative solution is required for SAP authorization management.

If the scope of security authorization policy management is defined as an SAP domain, SAP software provides capabilities for authorization management integrated with the SAP tools and run time. However, if the enterprise requires a centralized security authorization policy management solution that spans across SAP and non-SAP enterprise applications, a different approach is required.

Tivoli Security Policy Manager can use existing identity stores, such as IBM Security Identity Manager and IBM Security Access Manager for Web data stores, to form the basis of policy groupings. A policy-authoring component exists that can take non-technical documentation and translate it into implementable rules. Also, a policy lifecycle capability exists, so that policy rules can be approved, implemented, and periodically reviewed.

SAP software has its own proprietary mechanism to define and enforce security policies, and does not support XACML as of July 2014. However, there are benefits of Tivoli Security Policy Manager in a scenario with an SAP environment.

Several benefits of integrating SAP environments with centralized enterprise policy management are as follows:

- Define and create a central policy for managing message protection for applications and middleware surrounding SAP environments, such as authentication proxy, token services or ESB.
- Integrate SAP systems into an enterprise SOA environment.
- Manage authorization policies for access to SAP environments.

Figure 11-9 shows a generalized use case scenario with Tivoli Security Policy Manager in an SOA environment with an SAP web services client.

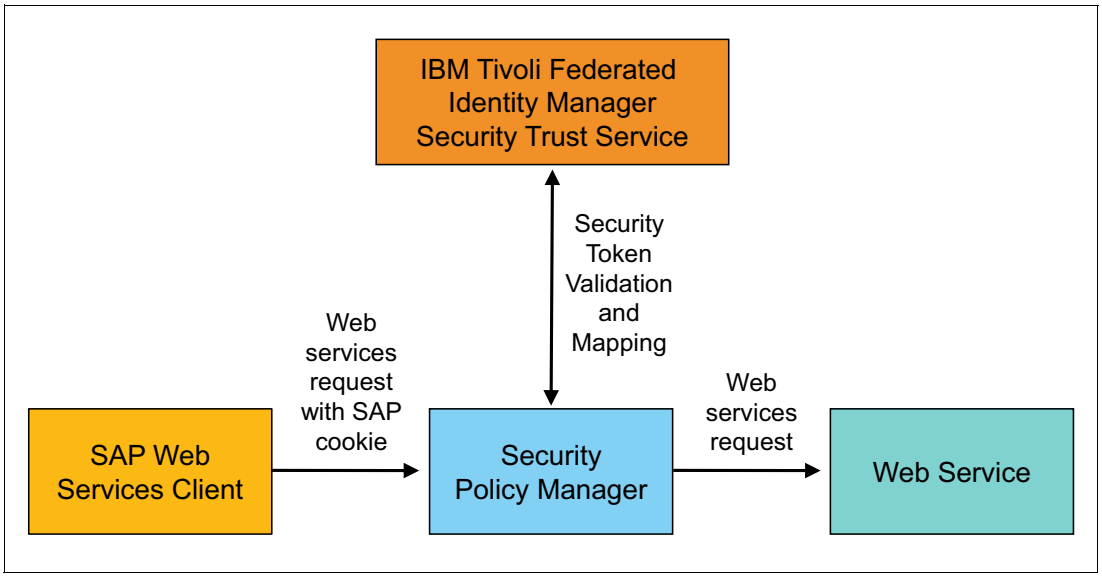

Figure 11-9 Security policy management in an SOA environment

11.6 Audit system scenarios

The objective of an SAP ERP audit system is to provide management with an independent assessment that relates to the effectiveness of the configuration and security of the enterprise's SAP ERP architecture.[2]

One of the responsibilities of an enterprise audit system is to ensure that all aspects of the security implementation comply with legal and corporate instructions, standards, and guidelines.

[2] Also see *Security, Audit and Control Features - SAP ERP, Audit/Assurance Programs and ICQs, Technical and Risk Management Reference Series, 3rd Edition*, ISACA, ISBN 978-1-60420-115-4.

The focus is on the primary corporate requirements, such as those listed in Table 11-7.

Table 11-7 Corporate audit requirements and examples

Audit requirement	Example
IT security standards	► User IDs ► Required system and client parameter settings ► Security and system administrative authority (prohibited and enabled) ► Security audit log settings ► Regular health checking requirements ► Access authorizations, including approval requirements ► Activity logging and application audit trails
Separation of duties standard	► Aggregate SOD ► Mitigations ► Role documentation requirements ► Use of GRC SOD tools
Application systems control and auditability, and responsibility for application owners and business process owners, addressing requirements for certification and compliance reviews and self-certifications	► Architectural overview ► SOD evaluation ► Application system management controls. including change management, problem management, access management, User ID creation and Access Assignment, User ID verification, and revalidation and compliance with corporate instructions and standards ► Data protection and privacy ► Test documentation requirements ► Business process owner acceptance ► Audit trails ► Risk evaluations
Classification and control of corporate information	► Classification of data ► Requirements for treatment of confidential data
Guidance on treatment of personal information and sensitive personal information	N/A
SAP application audit requirements	► Custom executable code protection ► Testing documentation requirements ► Critical objects ► Role documentation components

For information about audit products that are available from IBM, see 11.9, “Audit products and solutions” on page 302.

11.6.1 Security monitoring and analytics scenario

SAP environments consist of many layers that require monitoring to accomplish these goals:

- Gain visibility into network, user, and application activity to provide organizations with intelligence into potential and existing threats across their entire network.
- Fulfill regulatory and compliance mandates that require log management and tracking user activity.

SAP security monitoring should provide visibility, compliance, and risk management across the many layers required to support the integrated enterprise SAP environments (see Table 11-8 on page 295).

Table 11-8 Layers of SAP security monitoring

Security monitoring layer	SAP coverage
Presentation layer	► Systems hosting GUI services to ensure that servers are secure ► Identity and access services to make sure that users accessing the GUI have proper authority
Network layer	Network traffic through the systems supporting SAP services to look for potential and existing threats
Application layer	SAP application audit logs to monitor SAP user activity
Host layer	Servers hosting SAP services to monitor privileged user activity
Database layer	Databases used by SAP applications to monitor data access and database configuration changes

The selection of an SAP security monitoring architecture depends on the required depth and information detail (for example, transaction information details), and the scope of the integration (such as SAP environment-specific or enterprise-wide scope). Table 9 shows an example of security monitoring alternatives and architectural decisions.

Table 11-9 Security monitoring architectural decision

Subject	SAP security monitoring
Architecture decision	Define components and deployment architecture for SAP security monitoring that is detailed enough to be SAP environment-specific, and generic enough for enterprise compliance reports.
Alternatives	► Real-time transaction analysis and reporting ► User authentication logging and analysis (user login logging and illegal act monitoring) ► User authorization report (roles, groups, and SOD) ► SAP system security configuration analysis and reporting
Decision criteria	► If general and company-wide user authorization and activity reporting is in focus, use reports from IBM Security Identity Manager and IBM Security Access Manager that include activities on the SAP system. ► For real-time analysis, use sensor technology to configure alerts in case of misuse, intrusion, and so on (such as IBM QRadar Security Intelligence Platform). ► For activity analysis and reporting on the SAP database, use a database audit and protection tool that is specific to SAP environments, such as IBM InfoSphere Guardium. ► To gather information from different systems and several SAP applications, and to correlate with other systems' information for enterprise compliance reports and auditing, use security information and event management (SIEM) software, such as IBM Security QRadar SIEM.
Decision example	► Correlate and analyze available information about user behavior. For SAP environments, this includes analyzing the SAP security audit log and system log files, or monitoring privileged SAP users. ► Collect and accumulate additional information to analyze how SAP systems are accessed, or to detect potential misuse, such as network traffic, database activity, and so on.

Subject	SAP security monitoring
Architecture	Use existing system information and logging, and place additional sensors and hooks where detailed security information is required (after weighing risks and performance effect), and use a correlation engine to interpret SAP security-related activities across the enterprise.

11.6.2 Source code analysis scenario

SAP application customers run custom-developed ABAP programs in their IT infrastructure, and the presence of vulnerabilities in these systems introduces several risks. Every custom development adds specific functionality to applications, often without any testing for risks or vulnerabilities.

The traditional time for a security audit is just before an application goes into production. A more prudent and cost-effective approach is to find the defect in development, at build time, or in QA.

Many application security teams are moving beyond annual audits and preproduction testing to adopt the leading practices of application risk management. The IBM Security AppScan portfolio includes capabilities to measure, monitor, and reduce the enterprise risk introduced by application vulnerabilities.

With the addition of Virtual Forge CodeProfiler for IBM Security AppScan source software, ABAP applications can be included in the enterprise-wide view of application risk. Security executives and auditors can benefit from more than 40 compliance reports that are built into the software, and trending analysis that is useful for measuring and driving the reduction of application risk.

Virtual Forge CodeProfiler for IBM Security AppScan source software is a static analysis security testing solution that helps to identify vulnerabilities in ABAP source code, review data flows, and identify threat exposures in SAP applications (see Figure 11-10).

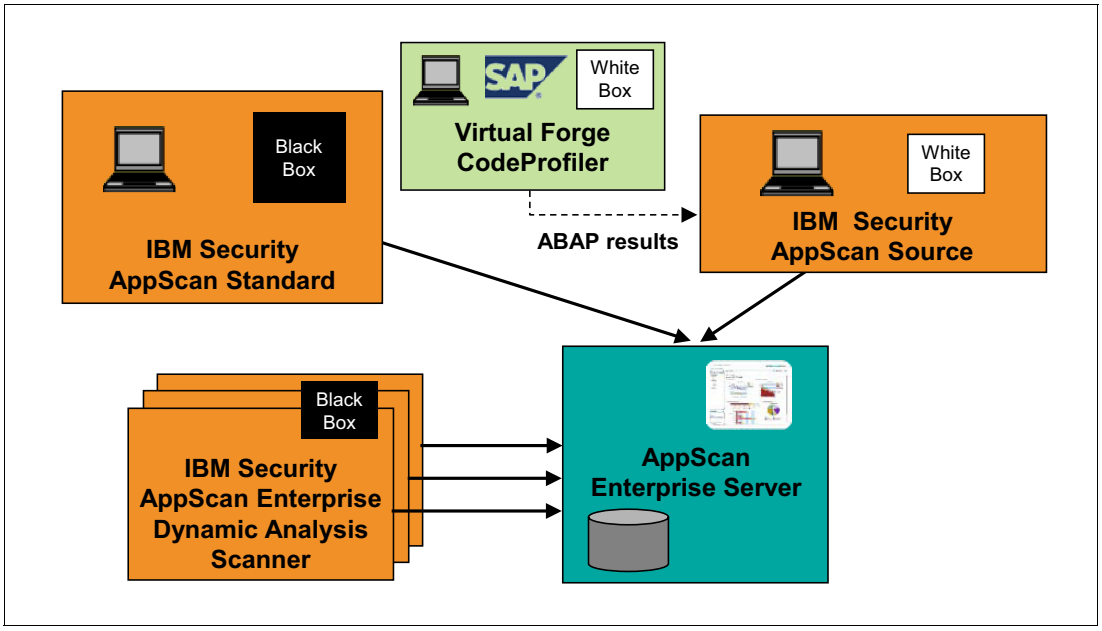

Figure 11-10 IBM Security AppScan and Virtual Forge CodeProfiler

11.7 Identity management products and solutions

This section describes the main products and solutions in the IBM identity management portfolio that can be used when implementing SAP system security scenarios.

11.7.1 IBM Security Identity Manager

IBM Security Identity Manager (formerly known as Tivoli Identity Manager) is a policy-based identity and access governance solution that helps automate lifecycle management of user roles, identities, and access rights. Through the use of roles, accounts, and access permissions, IBM Security Identity Manager helps automate the creation, modification, and termination of user privileges throughout the entire user lifecycle.

IBM Security Identity Manager enables centralized management and administration of users within the IT environment of an enterprise. The following list describes management and administration functions that are provided by IBM Security Identity Manager:

- User account provisioning
- User account password management
- Account request approval workflows
- Account recertification
- User access role and group membership management

A large inventory of adapter components enables IBM Security Identity Manager to manage separate, distinct IT applications and resources within heterogeneous environments. Adapters are deployed as separate installable units within the infrastructure.

IBM Security Identity Manager supports the IBM Security Identity Manager Adapter for SAP NetWeaver, which enables account management for the SAP NetWeaver Application Server ABAP server stack. This adapter enables administration and provisioning of user accounts between IBM Security Identity Manager and SAP NetWeaver ABAP server. It also includes optional integration components, enabling integration between IBM Security Identity Manager and SAP GRC Access Control.

11.7.2 IBM Security Directory Integrator

IBM Security Directory Integrator integrates and synchronizes generic and identity data in a variety of system stores, such as files, message queues, web services, directories, databases, collaborative systems, and applications used for HR, CRM, ERP, and other corporate applications. It can help organizations build an authoritative data infrastructure, enabling consistent data across multiple identity or generic data resources.

IBM Security Directory Integrator performs the following tasks:

- Transforms, moves, and synchronizes generic and identity data in heterogeneous directories, databases, files, collaborative systems, and applications, with real-time automated updates to the authoritative data source.
- Helps enhance the security, accuracy, and integrity of generic and user identity data, while facilitating data migration, transformation to other file formats, and synchronization between two or more systems.

The preferred framework for new IBM Security Identity Manager adapter development is based on IBM Security Directory Integrator. Adapter implementations are embodied by IBM Security Directory Integrator AssemblyLines. AssemblyLines ideally use one or more IBM Security Directory Integrator connectors or function components to facilitate target resource interfacing with additional Java or JavaScript processing components.

IBM Security Directory Integrator includes a set of connectors for integration with SAP systems. These connectors are combined within the IBM Security Directory Integrator Component Suite for SAP NetWeaver Application Server ABAP.

Installing IBM Security Directory Integrator also installs the Component Suite for SAP ABAP Application Server. However, to complete the install of the Component Suite, an additional component must be added on the target machine if it does not already exist, which is the SAP Java connector2 (SAP JCo2).

The IBM Security Directory Integrator Component Suite for SAP ABAP Application Server includes the following components:

- Function Component for SAP NetWeaver Application Server ABAP
- User Registry Connector for SAP NetWeaver Application Server ABAP
- Human Resources/Business Object Repository Connector for SAP NetWeaver Application Server ABAP

11.7.3 IBM Security Directory Server

IBM Security Directory Server software provides a reliable platform for enterprise security initiatives. This enterprise identity management software from IBM uses LDAP to provide a trusted identity data infrastructure for authentication.

IBM Security Directory Server includes the following capabilities:

- Provides identity management for companies that want to deploy a robust and scalable identity infrastructure
- Uses LDAP identity infrastructure software, and meets LDAP v3 industry compliance standards
- Enhances proxy server capabilities with flow control for managing requests, paging search results for single and multiple partitions, and a smart fail-back mechanism to restore servers safely
- Maintains high availability with master and subordinate, and peer-to-peer, replication capabilities, and scheduled online or offline backup and remote restore
- Supports virtual list views so that you can scroll forward or backward through entries in a large sorted data set, and can record deleted entries

IBM Security Directory Server is an SAP-certified product for integration with SAP using the SAP standard interface BC-LDAP-USR. This integration includes interoperability with SAP NetWeaver Application Server ABAP and Java. As a result, IBM Security Directory Server can be used as the identity store for synchronization with the AS ABAP (application server for ABAP) user repository, and the persistence user store (user data source) for the AS Java User Management Engine.

11.8 Access management products and solutions

Managing access across IT resources, in most cases, means providing technologies that enable you to establish an instance to handle central authentication and SSO to enterprise systems and applications. IBM takes a divide-and-conquer approach to this challenge, and addresses different domains with different technologies. Three domains of authentication and SSO are as follows:

- Web environments
- Desktop environments
- Federated environments

Each of the three computing models has SSO requirements, and IBM applies different technologies to meet each of the SSO requirements. These technologies are implemented in the following products:

- IBM Security Access Manager for Enterprise Single Sign-On
- IBM Security Access Manager for Web
- Tivoli Federated Identity Manager

11.8.1 IBM Security Access Manager for Enterprise Single Sign-On

This product addresses the enterprise SSO problem by deploying an agent on the desktop that intercepts authentication requests by applications, and automatically completes the login data with credentials stored on the local machine.

IBM Security Access Manager for Enterprise Single Sign-On AccessProfiles is composed of structured XML files that enable SSO automation for applications. These profiles are created using the AccessStudio component, and they are essentially used to automatically capture and inject application credentials in a user's "wallet."

IBM Security Access Manager for Enterprise Single Sign-On integration with SAP solutions includes support for IBM Security Access Manager for Enterprise Single Sign-On AccessProfile for SAP GUI for Windows.

11.8.2 IBM Security Access Manager for Web

This product addresses the web SSO problem by placing a reverse web proxy in front of the enterprise web applications. IBM Security Access Manager user accounts are stored in an enterprise directory, and users only need to authenticate to the IBM Security Access Manager server to access all of the existing web applications configured behind the reverse proxy.

IBM provides the following integrations between IBM Security Access Manager for Web and SAP solutions:

- IBM Security Access Manager for Web integration with SAP NetWeaver Application Server ABAP
- IBM Security Access Manager for Web integration with SAP NetWeaver Application Server Java
- IBM Security Access Manager for Web integration with SAP NetWeaver Application Server Java Enterprise Portal Core
- SSO for SAP NetWeaver Application Server ABAP with IBM Security Access Manager for Web in conjunction with SAP NetWeaver Application Server Java

11.8.3 IBM Tivoli Federated Identity Manager

This product addresses the federated SSO problem by implementing industry-standard federated SSO protocols, including SAML, Liberty ID FF, WS-Federation, Information Card, Open ID, and OAuth. It supports arbitrary identity transformations based on Extensible Stylesheet Language Transformation (XSLT), IBM Security Directory Integrator AssemblyLines, or custom Java programming, so that credentials can be converted to a format compatible with the local environment.

Tivoli Federated Identity Manager provides access management and SSO to federated SAP environments and SAML-enabled SAP applications.

Restriction: Tivoli Federated Identity Manager supports SAML; however, IBM does not provide a formally supported integration based on SAML for SAP software. The integration is technically feasible, but IBM is not committed to accept support APARs if problems occur.

The article *Integrating IBM Federated Identity Manager 6.2.2 with SAP Login Tickets* shows an example of an unsupported integration. Customers or their service provider are responsible for implementing and maintaining Login Ticket-based integration. The article is available at the following location:

`http://ibm.co/1l2xlpV`

Tivoli Federated Identity Manager integration with SAP NetWeaver Application Server can be enabled based on the following unsupported integrations:

- SSO using SAML Browser Artifact (SAML integration)
- STS for the SAP Login Ticket (SAP Token Trust Module)

Tivoli Federated Identity Manager STS is a WS-Trust compliant service that enables users to validate, exchange, and issue tokens. Within the service is a set of module chains called *trust chains*. Having many trust chains within the service is possible. To determine which chain to start, it looks at the `AppliesTo` and `Issuer` elements (among others) in the request. Each trust chain consists of one or more modules that can either validate, map, issue, or exchange tokens.

By developing a custom STS mapping module, the IBM Tivoli Federated Identity Manager STS can validate identity tokens containing SAP login tickets. It is a trust module that can be used in a trust chain to validate an SAP user identity issued in an SAP Login Ticket by an SAP system.

Figure 11-11 provides an overview of the validation process when used in an STS trust chain.

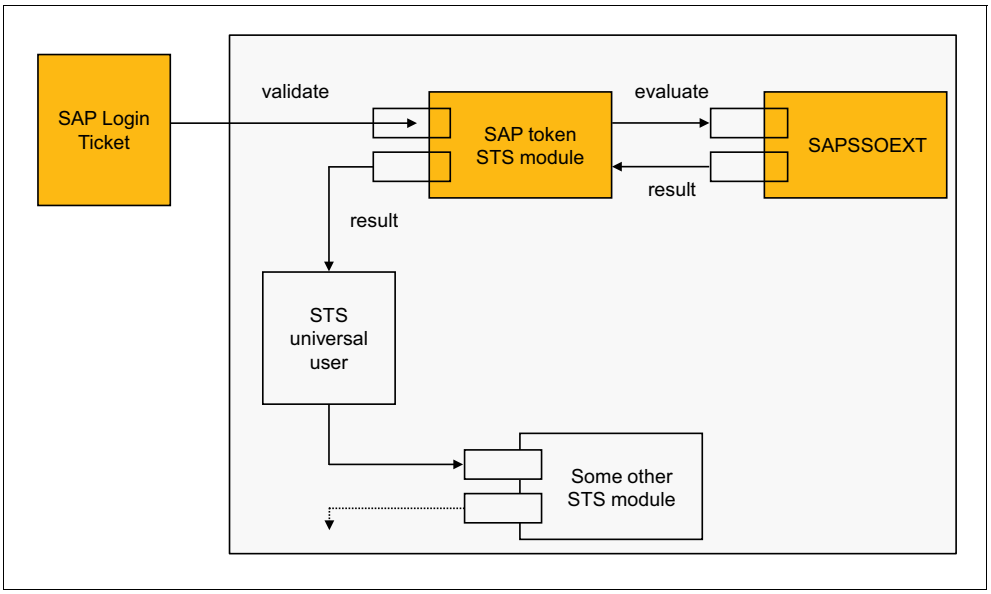

Figure 11-11 STS trust chain validation process

The technique can be used in conjunction with WebSphere DataPower XML firewall with the WebSphere DataPower appliance placed inline and to act as a proxy. During processing, it performs a WS-Trust call to the Tivoli Federated Identity Manager STS to exchange the SAP identity token, which was sent as a cookie, for a new token. This new token is placed in the web service request as a WS-Security header, and forwarded to the service. Figure 11-12 illustrates a high-level view of the solution.

This solution architecture delivers the following benefits:

- Integrates SAP systems into an SOA environment by converting requests into open standards-compatible messages.
- Provides extensible design that reacts to future requirements.
- Requires minimal changes to existing infrastructure.
- Propagates identity from end-to-end, enabling authorization and auditing at every node.

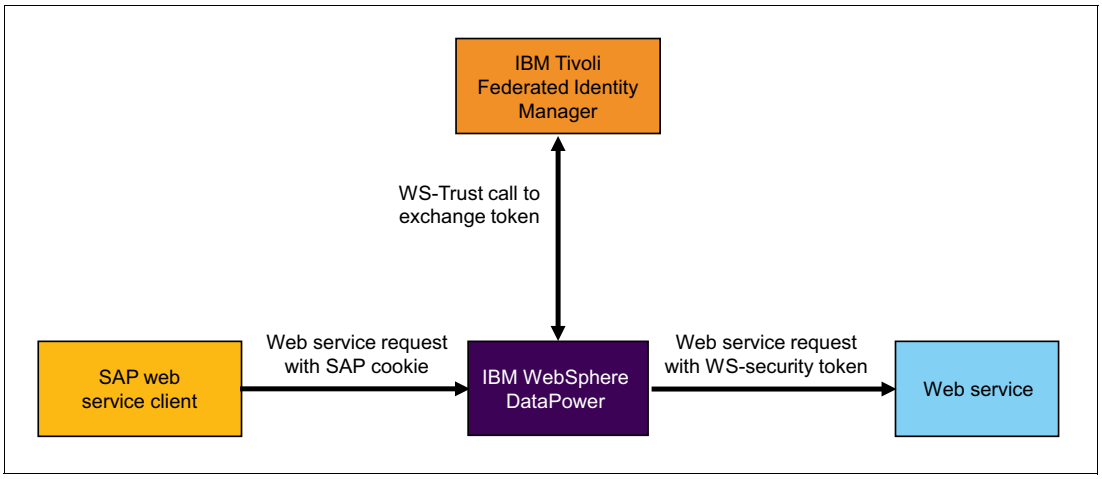

Figure 11-12 Federating the SAP login ticket with Tivoli Federated Identity Manager and WebSphere DataPower

11.9 Audit products and solutions

This section describes the main IBM products and solutions that can be used to add extensive audit and logging capabilities to SAP systems security.

11.9.1 IBM InfoSphere Guardium

Guardium integrates with SAP solutions to identify an individual user from a pooled database connection. It provides predefined policies and reports for SAP systems to comply with Payment Card Industry Data Security Standard (PCI DSS), Sarbanes-Oxley (SOX), and ISO27001. In addition, IBM InfoSphere Guardium Encryption Expert can encrypt all SAP database files and provide compliance without making any changes to the applications or databases.

With Guardium, detailed information is available about the SAP solution users and their activity. This information reaches beyond SAP transaction logs, and makes it easier to detect fraud and unauthorized activity. Guardium integrates with SAP solutions and SAP databases with no application changes required.

11.9.2 IBM Security QRadar Log Manager

IBM Security QRadar Log Manager provides enterprise log management for compliance, and forensics for SAP environments. It provides comprehensive, high-performance, and tamper-proof log management for organizations looking to collect, archive, secure, and analyze large volumes of network and security event logs for security, auditing, and reporting requirements in the environment running SAP. The SAP integration features include the following capabilities:

- Management of network and security events in the network that the SAP environment is running on
- Management of host logs on operating systems underlying each SAP server
- Management of the SAP ERP application logs
- Management of the database activity on databases used by SAP services
- Tamper-proof data archive of security event data from the systems and networks used in the SAP environment
- Threshold-based correlation and alerts based on data from the SAP environment
- Reporting templates support compliance mandates:
  - Control Objectives for Information and Related Technology (COBIT)
  - SOX
  - Gramm–Leach–Bliley Act (GLBA)
  - North American Electric Reliability Corporation (NERC)
  - Federal Information Security Management Act (FISMA)
  - PCI
  - Health Insurance Portability and Accountability Act (HIPAA)
  - UK Good Practice Guide/Group Control System (GPG/GCS) 13
  - ISO 27001

11.9.3 IBM Security QRadar Risk Manager

IBM Security QRadar Risk Manager provides Enterprise Risk Management across SAP environments. On systems and networks running SAP services, it provides advanced intelligence risk management capabilities in the areas of policy monitoring, device configuration and topology, network and security event modeling and simulation, and advanced network visualization. The SAP integration features include the following capabilities:

- Monitoring of network security configuration and sending of alerts when risky or out of compliance configuration is found in systems or the network where SAP is running.
- Read-to-use policy templates help assess risk across multiple regulatory mandates.
- Enables network activity to monitor and support forensic search, visualization tools, and out-of-policy flagging using the data collected from the SAP environment.
- Analysis of network and security events to assess policy effectiveness and determine most up-to-date configuration changes on systems and the network used to provide SAP services.
- Integration of network topology with vulnerability scan results to better assess which SAP systems are most vulnerable.
- Vulnerability modeling, simulation, and visualization.

11.9.4 IBM Security QRadar SIEM

IBM Security QRadar SIEM provides enterprise visibility and threat intelligence for SAP environments. It provides capabilities for real-time analysis, threat detection, and incident management with sophisticated correlation using data from the SAP environment. The following list describes the SAP integration features:

- Complete QRadar Log Manager functionality.
- Real-time analysis, threat detection, and incident management based on data from the SAP environment.
- Sophisticated correlation, incorporating diverse data sets collected from the SAP environment (event, flow, asset, vulnerability, and external intelligence).
- Managing flows for full network behavior analysis.
- Asset profile creation and management for all assets in the SAP environment.
- Incident management for problems found in the SAP environment.

Integration of IBM Security QRadar SIEM with IT-Cube agileSI provides SAP security monitoring with the combined capabilities of IBM Security QRadar SIEM and IT-Cube agileSI. IBM Security QRadar provides capabilities for real-time analysis, threat detection, and incident management using sophisticated correlation with data from the SAP environment.

The solution analyzes several SAP parameters, system information, and data stored in log files, tables, change documents, and so on. Table 11-10 lists potential events and data, and sample scenarios.

Table 11-10 SAP security events and use cases

Events and data	Sample use cases
SAP Security Audit Log Subset of security events in SAP systems, such as (failed) logins, transaction starts, and so on.	Brute force login: ► User created, deleted, locked, or unlocked ► Password changes ► Execution of reports
SAP basis log for availability, error tracking, security, and so on	► Debugging ► Execution of operating system (OS) commands
SAP system configuration	Password policy checks: ► SAP Gateway check ► Encryption of communication (SNC status)
Data stored in tables	System and client change settings: ► SSO, logon tickets ► RFC configuration ► Any data stored in any table
Monitor availability	Check availability of SAP systems
Communication with external programs	Monitor *denied* external calls
Authorization data	SOD checks
Changes to data stored in tables	Monitor critical tables (master data and conditions of purchase)

11.9.5 IBM Security AppScan

IBM Security AppScan with Virtual Forge CodeProfiler provides vulnerability scanning for SAP ABAP and enterprise applications (Static Analysis). Integrating SAP ABAP Source Code scanning with enterprise application source code scanning offers the following benefits:

- Helps ensure the security and compliance of web, mobile, and SAP ABAP applications by scanning the source code for application vulnerabilities
- Identifies emerging application vulnerabilities
- Helps prevent exploitation by hackers
- Early detection and resolution of web application vulnerabilities decrease the risk of attack
- Saves valuable development rework, time, and resources
- Lowers the total cost-of-ownership because vulnerabilities can be corrected early in the product development cycle, before becoming security risks
- Static (*white box*) analysis of client-side issues, such as Document Object Model (DOM)-based cross-site scripting (XSS) and code injection, to secure SAP and non-SAP web portals, applications, and SOA middleware

IBM Security AppScan static analysis provides the following capabilities:

- A Static Application Security Testing (SAST) solution (source code scanning) for various languages, including Java, for both web and mobile applications
- Integrated security testing as part of the application development process, including at build time
- Automated aggregation of results with other analysis to provide a single integrated repository for reporting and analysis

CodeProfiler for SAP ABAP applications adds the following capabilities:

- Source scanning of SAP ABAP applications
- Integrates with IBM Security AppScan
- Partnership with Virtual Forge GmbH
- Enforce automates code analysis for security defects, standards, and service level agreements (SLAs)
- Regains control of ABAP code
- SAP integration features

IBM Security AppScan dynamic analysis provides vulnerability scanning for enterprise applications (dynamic analysis). Integrating dynamic analysis scanning with SAP ABAP and enterprise application source code scanning offers the following benefits:

- Automated dynamic (*black box*) testing of web vulnerabilities and web services, Web 2.0, and rich Internet applications (JavaScript, Ajax, and Flash).
- Secures SAP and non-SAP web portals and applications and SOA middleware.
- Completes the customer solution provided by IBM Security AppScan static analysis and Virtual Forge CodeProfiler.
- Helps ensure the security and compliance of web applications, mobile applications, and SAP ABAP applications throughout the software development lifecycle.
- Combines results of static and dynamic scanning.
- Desktop deployment provides on-demand Dynamic Application Security Testing (DAST).
- Collaborative deployment provides server-based, parallel dynamic scanning.
- Simulates attacks to expose vulnerabilities.
- Provides application security scanning and centralized reporting.

11.10 References

For more detailed information about the topics in this chapter see the following publications:

- *Using the IBM Security Framework and IBM Security Blueprint to Realize Business-Driven Security,* SG24-8100
- *Integrating IBM Security and SAP Solutions*, SG24-8015

12

Systems management for SAP

Availability, supportability, and maintainability of systems of record should always be a primary concern for organizations. The criticality of the types of business processes typically enabled by an enterprise SAP transformation, however, makes these concerns paramount, forcing a key architectural decision point (DP) early in project planning. Organizations must ask themselves the following typical questions:

- How will systems management be addressed?
- Will SAP be considered as its own category and managed using its own toolset, while using other tools to manage the rest of the technical solutions spanning the enterprise?
- Will heterogeneous applications and infrastructure be managed by integrating them into SAP's system management tools? Or will SAP's system management tools be integrated into the broader, enterprise-wide systems management solution?

This chapter describes the goals of enabling business process availability management in SAP deployments through the integration of systems management tools across a heterogeneous environment. This chapter provides an overview of the IBM systems management architecture for SAP, which supports a business services management model, describes its constituent components, and explores the rationale for the choice of a particular component.

This chapter includes the following topics:

- 12.1, "Architectural goals" on page 308
- 12.2, "Business process availability management overview" on page 309
- 12.3, "Systems management reference architecture for SAP-centric solutions" on page 312
- 12.4, "Business process availability management for SAP-centric solutions" on page 315
- 12.5, "Summary" on page 320
- 12.6, "References" on page 321

© Copyright IBM Corp. 2014, 2015. All rights reserved.

12.1 Architectural goals

Large SAP deployments require an automated solution for managing business process availability, to deliver highly available processes to the business. A business process availability management solution takes into account the effect on business processes of an integrated dependency model between infrastructure, applications, and business processes.

This approach enables the actual management of business services, rather than the management of the individual underlying systems and applications as a proxy for the business processes that they support.

This section describes the main architectural goals to consider when designing systems management into SAP solutions.

12.1.1 Enable optimal availability and usability of complex business systems

The architectural goal for system management of an enterprise application built on SAP software is to consistently enable optimal availability and performance. Accounting for every component, every piece of infrastructure, and every application running in the data center is a daunting enough task.

However, if an organization merely ensures that they have a tool to provide a view of the availability of every bit of infrastructure, and every application, with no regard to how those subsets of the overall solution affect the business processes that are important to their users, they are left with a whole that is substantially less than the sum of its parts.

Enabling optimal availability and usability of the critical business processes of an enterprise, and all of its supporting technical bits and pieces, is best accomplished by implementing a business service management (BSM) model that integrates the myriad systems management tools with the support processes and teams needed to operate them. The implementation of such a BSM model is a non-trivial accomplishment, especially with the level of technical complexity that an SAP transformation can engender.

A business manager *does not care* that the system that supports accounts payables for all of Asia Pacific is up and running. That manager only cares whether the accounts payables users are able to access that system and complete their tasks in a timely fashion.

Traditional information technology (IT) systems management has used *system* availability as a proxy for *business process* availability, but the onward push toward consolidation of business systems, and the complexity of the relationships of the processes to the systems, make this approach increasingly infeasible.

12.1.2 Provide visibility to unplanned business process outages

If a subcomponent fails or becomes unresponsive, an effective systems management solution enables both the IT and business support teams and management to see, in real time, the effect of the problem on the business processes of the company. Is a subcomponent failure rendered moot by built-in architectural redundancy? Is it less consequential because it is affecting a process at an off-peak time?

For a subset of critical business processes, a real-time view of availability can be obtained through direct instrumentation or synthetic transaction monitoring (probing). However, those approaches are difficult to scale up to be able to cover the entirety of the business processes of an enterprise. Also, not combining business process-specific monitoring with the vast amount of monitoring available at the application and infrastructure levels of large systems is a wasted opportunity.

The ability to dynamically link a business process hierarchy with its dependent application and infrastructure relationships is crucial for minimizing problem determination time and aiding decision making. This approach, in turn, enables for more holistic and less siloed support practices, and more efficient use of support resources.

12.1.3 Enable historical view of business process availability

Another key milestone toward bridging the gap between what business managers want and what their IT service providers deliver is the ability to report service level agreements (SLAs) for individual business processes. After an enterprise has the ability to track the true availability of their core business processes, in real time, it is a small step from there to retain and mine that data to enable historical views of business process availability.

A historical view of business process outages enables for correlation of business process availability with capacity and non-critical fault trends. This information is of great help to an organization in understanding weaknesses in their architecture, or for prioritizing investments at the application and infrastructure level to maximize value to the business.

12.2 Business process availability management overview

This section provides an overview of common challenges and requirements encountered by many enterprise customers, and it describes the proposed solutions at a high level.

12.2.1 Complex IT solutions require multiple levels of systems management

Multiple levels of monitoring and management are required by enterprise customers to ensure availability of their critical business systems. An IT organization must be able to support and manage the servers (physical and virtual), operating systems, connections, and storage that compose the infrastructure.

It must be able to similarly look after the applications, including, but not limited to, multiple SAP products, middleware, third-party software, and custom-developed business systems that run on the infrastructure and enable the execution of critical business processes. Finally, the successful IT organization must be able to support and manage the availability, performance, and usability of the business processes that are run by its users, partners, or customers.

Absent a comprehensive business services management model, these IT management layers are viewed, managed, and operated independently. No linkage exists within the technical management and monitoring of Infrastructure, applications, or business processes, as shown in Figure 12-1.

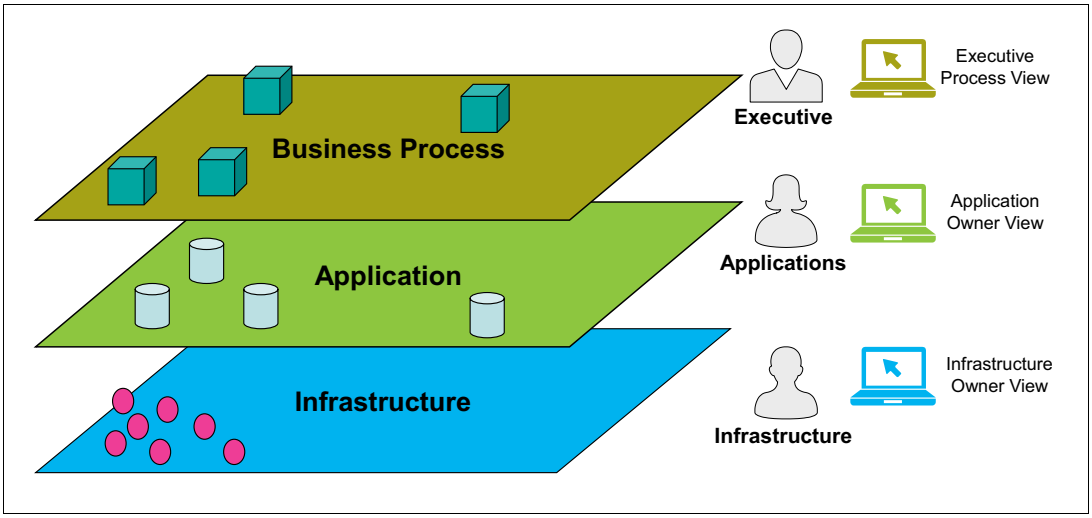

Figure 12-1 Independently managed, multi-layer IT services environment

12.2.2 Multiple systems management tools exist for each layer of solution

Management of the multiple levels of complex business solutions, including large SAP implementations, can be a daunting project because of the many systems management and monitoring options available for each of the layers of the technical environment. Often, multiple choices also exist to manage technologies and applications that are, themselves, only subcomponents of a particular management layer.

Some tools are specific in their scope, while others claim to “do it all.” Often overlaps in coverage exist, with more than one product covering part of the functional solution architecture. At other times gaps in coverage exist, meaning no tool is available to specifically manage or monitor part of the enterprise solution. The skills required for operating, managing, and maintaining these subsets of large systems can vary wildly, resulting in different teams using different management tools.

A key challenge for a successful IT services organization is in the selection of these tools and products, and the optimal integration of them to enable seamless business operations. This integration enables the required levels of service to be predictably delivered to the business.

These systems management tools and solutions, along with the underlying relationships between the various technologies connected throughout the enterprise, can be crucial in determining the organizational support structure. They have a great bearing on the fundamental ability of a company to manage change, respond to problems, and, ultimately, deliver the return on investment (ROI) underlying the entire project.

Ideally, supporting a large, complex, and mission-critical business system requires a comprehensive BSM model that, in turn, needs integrated operations. This integrated BSM model accounts for internal dependencies between individual assets, so each layer and service delivery tower must be integrated.

Figure 12-2 shows the wanted horizontal and vertical integration of system management capabilities that enables integrated operational support.

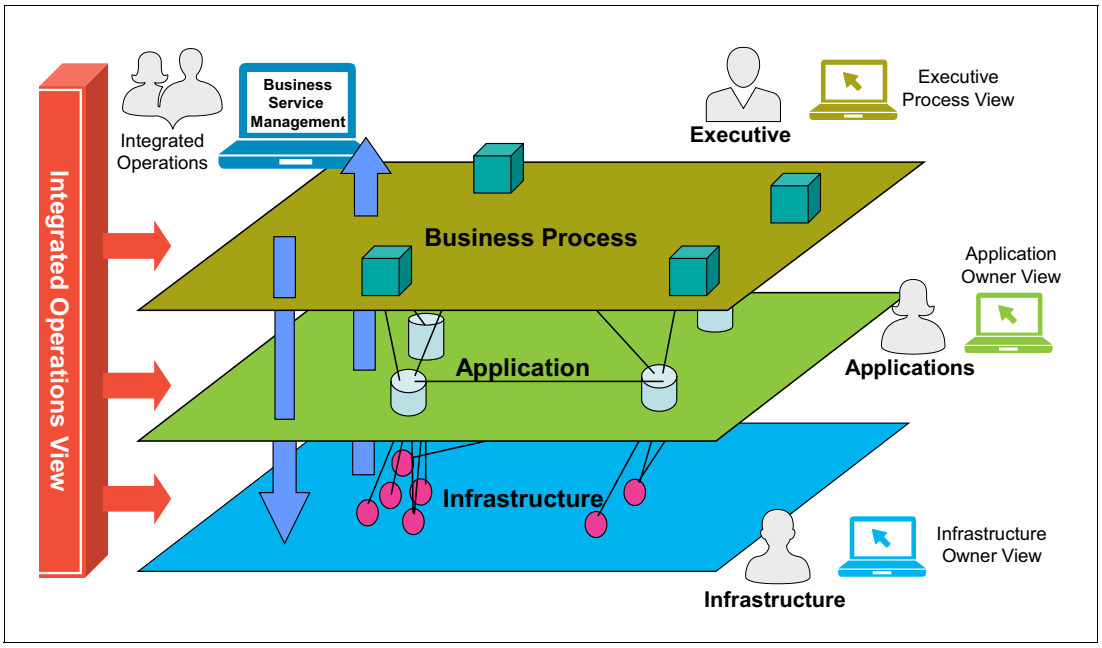

Figure 12-2 Integrated systems management within multi-layer IT services environment

As Figure 12-2 suggests, business processes, applications, and infrastructure are all interrelated. The execution of a specific business process can affect several applications, each of which in turn might depend upon multiple infrastructure components, such as servers, storage devices, and network switches.

12.2.3 Systems management considerations

Supportability, of course, is a key concern in the functional architecture of mission-critical systems, such as the SAP implementations described in this book. However, even as the SAP implementation project makes architectural and design decisions based upon functionality and cost, it should also strongly consider the effects of these decisions on supportability, availability, performance, and sustainability.

For example, a particular third-party business application might be the cheapest option for a specific business function, and might integrate with SAP well enough, but is it easily monitored and administered? Designing for supportability considers those types of factors alongside performance and functionality.

Another business application might have stellar native monitoring and management capabilities, but are those capabilities implemented in a way that enables for easy integration into the project's broader systems management solution? If not, does the application impose a significant cost upon the business, or introduce a major stability risk in the event that the integration difficulties cannot be mitigated?

If the project has to invest in heavy customization to achieve the wanted level of integration with its management and monitoring architecture, and then the monitored application changes, does all of the customization work need to be redone at a later time, for example, during a major version upgrade? Will the project experience a blackout in capability if a great monitoring tool is sourced from a company that goes out of business, or is bought by a competitor insistent upon bundling it with incompatible or redundant products?

12.3 Systems management reference architecture for SAP-centric solutions

This section reviews an example of a realistic, representative, SAP-centric deployment application architecture. Subsequently, it defines the corresponding monitoring architectures from a traditional systems management and a business process availability perspective.

12.3.1 Application architecture

Figure 12-3 depicts a drilled-down view of an SAP solution deployment architecture with multiple SAP applications, SAP partner applications, and IBM software components. The figure includes simplified relationships between the core SAP components (enterprise resource planning (ERP) Central Component (ECC), customer relationship management (CRM), business intelligence (BI), and so on), third-party applications, custom applications, portals, and middleware.

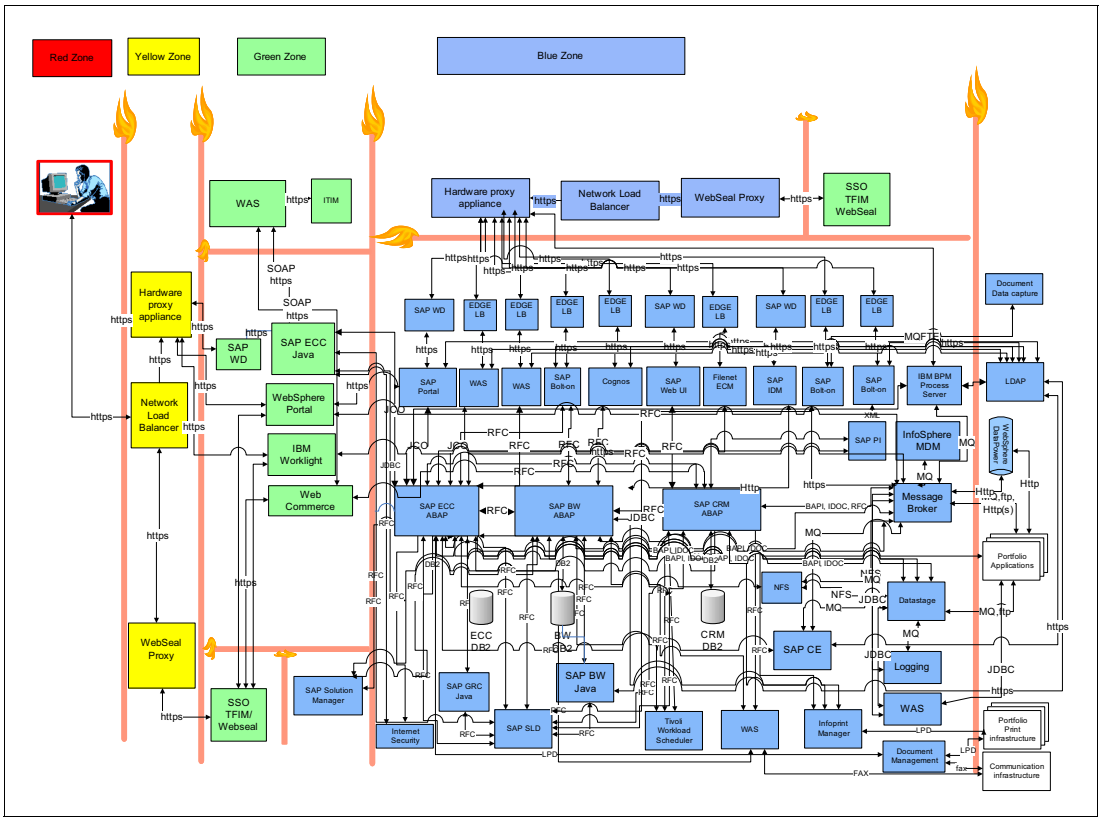

Figure 12-3 Partial detail of an SAP application architectural diagram

The network zones shown (red, yellow, green, and blue) represent differing levels of security and locations of components with regard to firewalls.

Each of the core applications (depicted by the blue rectangles in Figure 12-3) have several application subcomponents, and multiple physical and virtualized infrastructure components. All of these are combined in a multitude of interdependency relationships to enable hundreds of business processes. For example, the SAP ECC labeled *SAP ECC ABAP*, slightly to the left of the middle of Figure 12-3, has dozens of connections to other SAP business applications, customer business systems, messaging hubs, portals, and so on.

Viewed from an infrastructure perspective, that single blue SAP ECC rectangle is supported by multiple application servers, which span both virtual and physical machines. It connects to the database servers, which also have physical and virtual components, and which can be on a different technical platform from the application servers. The database and application servers are also connected to a storage infrastructure.

A given business process might have a hard dependency upon a web-facing portal, network authentication, and security infrastructure, a combination of SAP business applications, a third-party tax calculation system, and an internal messaging hub. In addition, all of those applications and components expand out to physical dependencies upon a lengthy list of virtual and physical system resources and their respective interconnections.

12.3.2 Infrastructure architecture

In parallel with the previous description of application architecture, the systems infrastructure has hard dependencies on physical and virtual machines and their constituent devices, from cards to cables. In an enterprise SAP-centric system, the infrastructure components are interconnected to enable for both steady state and failover operations.

The infrastructure components support multiple, overlapping applications. The machines, out of necessity, run on multiple operating systems. How is anyone supposed to understand, intricately, the interrelationships with these various technologies, and their interdependencies with the application architecture, let alone the combined relationships that exist between application, infrastructure, and a multitude of business processes?

12.3.3 Systems management architecture

The following systems management architecture (Figure 12-4) depicts IBM and SAP systems monitoring products, working together, combined to cover the application, infrastructure, and business process monitoring requirements of a large, SAP-centric business solution.

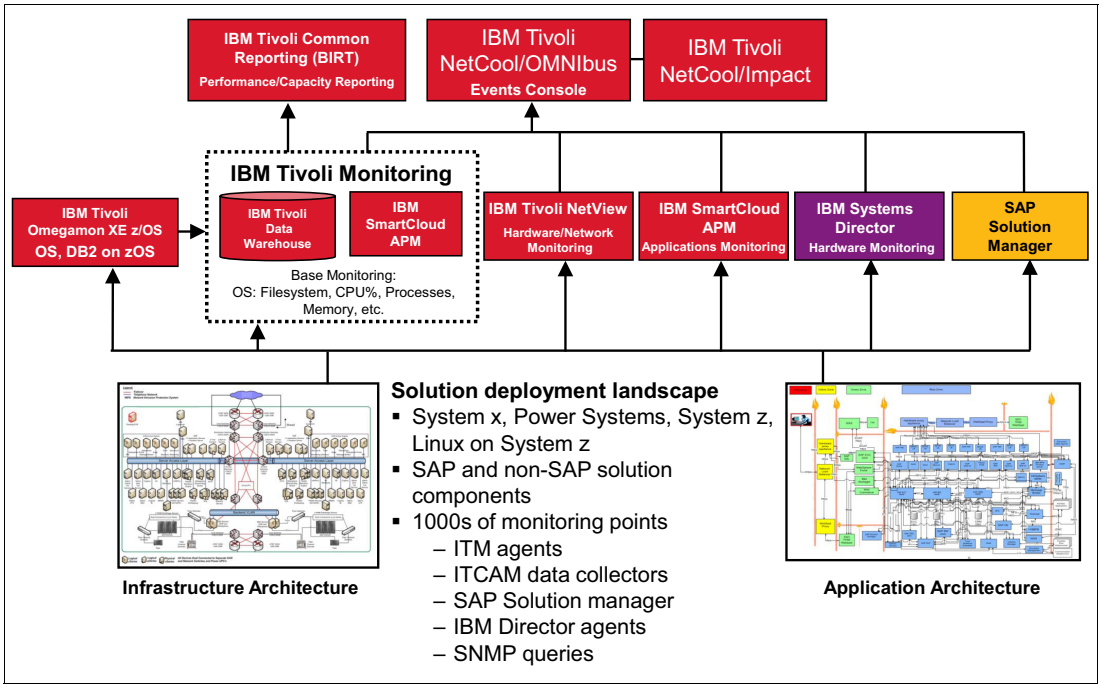

Figure 12-4 Systems management architecture

Be sure that the monitoring and management tools that compose the systems management architecture provide coverage over the totality of the technical solution. This extensive coverage, however, just provides the *potential* of visibility to problems and the ability to quickly diagnose and resolve them across the enterprise. Integrating these tools is crucial to maximizing ROI in the monitoring and management infrastructure and, by extension, the overall business systems.

Systems management integration decisions

Figure 12-4 on page 313 shows SAP Solution Manager as a part of this overall systems management architecture, but not as the central hub. Technically possible is to integrate infrastructure and non-SAP applications monitoring into SAP through SAP Solution Manager, but that approach requires extensive SAP Solution Manager customization.

For a smaller SAP implementation, this configuration effort might not be a significant factor in the decision. However, the greater the complexity of the physical infrastructure and non-SAP IT landscape, the greater the integration effort to funnel that technical information into SAP Solution Manager.

If a significant non-SAP systems and applications footprint exists and, especially, if those resources are already managed by other technical solutions (including the many IBM software products described in this book), consolidation of application and systems monitoring into SAP Solution Manager is a less practical approach.

For complex, heterogeneous IT environments, especially when considerable application and infrastructure monitoring is in place, it makes more sense to view SAP Solution Manager as a *source* of monitoring and management information, as opposed to the hub. Information from within SAP Solution Manager that provides significant insight into performance, health, and availability of the enterprise applications can be viewed in concert with other applications and infrastructure to make determinations about the overall business systems status.

This approach does not preclude the role of SAP Solution Manager as the monitoring and management hub for those support resources primarily concerned with SAP (for example, SAP Basis administrators, or Business Process support analysts). The ideal, however, is for the application and business process-specific information from SAP Solution Manager to be integrated into the broader enterprise systems management context.

Systems Management Products

The following systems management products (shown in Figure 12-4 on page 313) indicate their role in the systems management architecture:

- IBM Tivoli OMEGAMON® XE for z/OS. The enterprise infrastructure architecture shown in Figure 12-4 on page 313 is a three-tier SAP architecture, with the application servers running on IBM POWER7® processor-based servers and the database servers on IBM System z10® mainframe servers, running IBM z/OS. OMEGAMON XE for z/OS provides monitoring and management of IBM System z® and IBM zEnterprise solutions.
- IBM SmartCloud® Application Performance Management. IBM SmartCloud Application Performance Management (APM) is a solution that includes application discovery, user experience monitoring, transaction tracing, in-depth diagnostics, and data analytics and reporting in a single package. IBM SmartCloud APM includes IBM Tivoli Composite Application Manager for Applications, IBM Tivoli Composite Application Manager for Microsoft Applications, and IBM Tivoli Composite Application Manager for Transactions.

- IBM SmartCloud Monitoring. This is the core of the architecture described in 12.3.3, “Systems management architecture” on page 313. It includes IBM Tivoli Monitoring and IBM Tivoli Monitoring for Virtual Environments. IBM SmartCloud Monitoring provides visibility into systems infrastructure, including both virtual and physical environments. It also provides for monitoring of heterogeneous environments.
- IBM Tivoli Composite Application Manager Family. Multiple Tivoli Composite Application Manager products that collectively provide for detailed monitoring of applications spanning multiple components, in addition to the ability to delve deep into infrastructure subsystems to identify bottlenecks, flag inefficiencies, and determine root cause for application problems.
- IBM Tivoli NetView® for z/OS. Provides management, automation, and monitoring of z/OS networks.
- IBM Tivoli Network Manager IP Edition. Provides real-time management, monitoring, discovery, topology visualization, and root cause analysis for layer 2 and 3 networks including Internet Protocol (IP), Ethernet, and Multiprotocol Label Switching (MPLS).
- SAP Solution Manager. SAP Solution Manager is an SAP product for managing all aspects of SAP operations, including monitoring, SLA measurements, business process performance, technical and functional monitoring, and key performance indicator (KPI) reporting.
- IBM Tivoli Business Service Manager. Integrated dashboard capable of providing operational and business audiences with the service visibility to effectively manage real-time service health and business activity, including automated service modeling, service impact analysis, root cause analysis, and tracking of KPIs and SLAs.
- **I**BM Tivoli Netcool/OMNIbus. Provides consolidated event management, correlation, and forwarding for multiple monitoring sources.
- IBM Systems Director. Provides tools for discovery, inventory, status, configuration, system health, resource monitoring, system updates, and event management for IBM hardware.
- IBM Tivoli Data Warehouse. It is an embedded technology that provides a repository for data from IBM Tivoli Monitoring, IBM SmartCloud Application Performance Management, and other products, enabling for analysis of health, performance, and availability data within the managed infrastructure. Tivoli Data Warehouse is included in Tivoli Monitoring and IBM SmartCloud APM.

12.4 Business process availability management for SAP-centric solutions

Multiple IBM products used in supporting a large SAP-centric enterprise business application can be combined in a particular systems management architecture to provide business process availability management (BPAM).

This combination of products and capabilities, as shown in Figure 12-5 on page 316, enables contextual integration of business process, application, infrastructure, and systems management architectures. In that way, this BPAM architecture enables an enterprise to transition from an IT systems management (ITSM) to a BSM model of operations.

12.4.1 Business process availability management architecture

The BPAM architecture combines the systems management architecture from Figure 12-4 on page 313 with additional products to provide an integrated, business process-centric view of enterprise system health and availability, for both real time and historical (reporting) purposes. These additional products include, principally, IBM Tivoli Application Dependency Discovery Manager and Tivoli Business Service Manager, in combination with the Business Process discovery library adapter (DLA) and, optionally, the Tivoli Monitoring DLA.

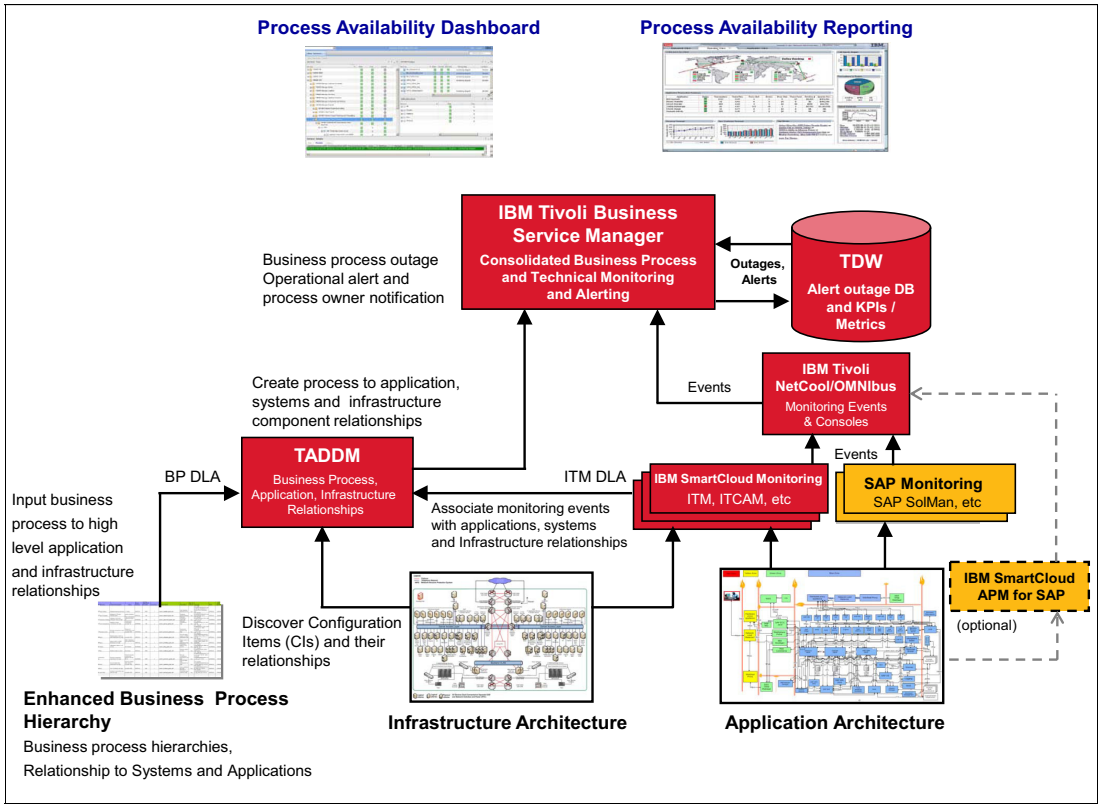

Figure 12-5 Business process availability management architecture

This solution consumes a customer-created representation of their business process relationships, called the *enhanced business process hierarchy* (BPH). The roles of these components in the BPAM architecture are explained in this section.

The following list describes the BPAM architecture functional flow and component descriptions (starting from the lower left corner of Figure 12-5):

- Enhanced Business Process Hierarchy. The enhanced BPH consists of a set of five comma-separated value (`.csv`) files that provide business process data in a format specified by the Business Process Discovery Library Adapter (BP DLA). Using the structured format of these `.csv` files, the customer declares the detailed interrelationships between business processes and their peer, parent, and child processes, in addition to documenting, for each process, their high-level application and system interdependencies.
- Business Process Discovery Library Adapter. The BP DLA consumes the enhanced BPH and translates these customer-declared relationships between business processes, applications, and systems into configuration items (CIs).

This is done in a database format recognizable by Tivoli Application Dependency Discovery Manager. To receive information about the source code and detailed instructions for the

Business Process DLA, contact Derek Jennings at dmjennin@us.ibm.com or Mathew Davis at mdavis5@us.ibm.com.

- Tivoli Application Dependency Discovery Manager. Tivoli Application Dependency Discovery Manager is the strategic data center discovery solution from IBM. It automatically creates and maintains application infrastructure maps based on sensors that discover application and system relationships, including configuration attributes and dependencies. These relationships are then used to create application and system topologies.

  In the BPAM architecture, the relationship and dependency data obtained by sensors from the application architecture and infrastructure architecture is augmented by the business process relationship data, sourced from the customer's enhanced BPH and loaded into Tivoli Application Dependency Discovery Manager using the Business Process DLA.

  The result is a Business Process←→Application←→Infrastructure integrated relationship model. Tivoli Application Dependency Discovery Manager takes the high-level dependencies declared by the customer between process and application, and application and system, and combines them with deep, low-level, detailed dependency information gained by interrogation of applications, systems, and subcomponents using sensors.

  For example, the BP DLA can declare that the sales order creation process has a dependency on SAP system YPE, which has a dependency on Application Server Group 001 and Database Server Group 002. Tivoli Application Dependency Discovery Manager, however, knows the precise hardware specifications of the constituent servers of Application Server Group 001.

  This information includes their physical and virtual machine names, their operating systems and other software release and patch numbers, and the connections those various servers have to storage. Because of the BP DLA, a direct line can be drawn between a business process and all of those lower-level dependencies.

- Monitoring. For the purposes of understanding the business process availability management architecture, remember that the lower right corner of Figure 12-5 is a simplified view of the systems management architecture shown in Figure 12-4 on page 313.

  IBM SmartCloud Monitoring, which consists of Tivoli Monitoring and Tivoli Composite Application Monitoring, is shown in Figure 12-5 along with SAP Application Monitoring (whether through SAP Solution Manager or IBM SmartCloud Application Performance Monitoring for SAP).

  However, many other monitoring products are available, and they cover various aspects of the application, infrastructure, and even business process architecture. Also, the exact components of this systems management architecture vary based on the types of applications and infrastructure, and the monitoring choices, for a given customer's SAP-centric enterprise system.

  For BPAM, the consolidated output of the monitoring infrastructure takes two paths:

  - Toward Tivoli Netcool/OMNIbus
  - Through the Tivoli Monitoring DLA and into Tivoli Application Dependency Discovery Manager

- IBM Tivoli Netcool/OMNIbus. Multiple sources of monitoring data feed into the Netcool/OMNIbus Event Manager, which enables for event correlation, consolidation, and forwarding. Although various monitoring events within an enterprise can go to different portals or dashboards visible to specialized support teams, Netcool/OMNIbus enables for monitoring event data to also be interrogated, filtered, and passed along based upon

configurable criteria. For BPAM, relevant monitoring event data is forwarded to Tivoli Business Service Manager.

- Tivoli Monitoring DLA. The Tivoli Monitoring DLA can be used to enhance the Tivoli Application Dependency Discovery Manager database of business process, application, and infrastructure interdependencies, with information about which monitoring exists for the various solution components.

  This is useful for reporting on monitoring coverage across the enterprise solution, enabling for identification of application and infrastructure components in need of monitoring. The Tivoli Monitoring DLA is available with the Tivoli Application Discovery Dependency product.

- Tivoli Business Service Manager. Tivoli Business Service Manager, in the context of BPAM, is a meta-portal. It is the place where the declared and discovered interdependencies can be displayed. The process/application/infrastructure dependencies are loaded from Tivoli Application Dependency Discovery Manager into Tivoli Business Service Manager, where they are used to build service trees.

  From these service trees in Tivoli Business Service Manager, an operator, manager, or technical specialist can drill down from a high-level process to a low-level one, and, from any level, identify the associated high and low-level application and system details, including real-time red/yellow/green status.

  Likewise, the physical or application infrastructure can be traversed, up or down. This enables outages or alerts to be viewed at the highest or lowest levels necessary to understand the potential scope of a problem. The health status of the items shown in Tivoli Business Service Manager is triggered by receipt of the events forwarded from Netcool/OMNIbus.

  Rather than configure every eventuality, rules can be set up so that inbound events will *light up* the appropriate node of the service tree when the ID of the affected components corresponds to the pre-configured contents of the process/application/infrastructure model imported from Tivoli Application Dependency Discovery Manager. This approach enables a comprehensive, yet flexible, solution.

12.4.2 Business Process DLA overview

The Business Process DLA is the key to enabling the business process availability management solution described in this chapter. Even when an organization realizes that a comprehensive mapping of their business process ←→application←→ infrastructure dependencies is of vital import, how can this be accomplished? The following list include typical questions organizations ask themselves in this context:

- Does our business process architect need to be an expert in monitoring tool configuration?
- Does our monitoring architect need to become an expert in multiple aspects of the company's business?

Each approach involves valuable resources spending time attempting to understand significantly different areas of expertise. The creation of the BP DLA resolves this dilemma, and it enables savings in enablement, administration, and maintenance of a comprehensive enterprise system monitoring and management solution.

The BP DLA enables a business process subject matter expert to work with an application architect to create a high-level model of business process relationships with other business processes, business process relationships with applications, and applications relationships with infrastructure. In the structured source (`.csv`) files that feed the BP DLA, the business

processes (BPs) are associated with system identifiers (SIDs), which can either correspond to actual SAP SIDs, or be created for the sake of reference to solution components.

These BPs and SIDs are correlated by the Business Process DLA with Tivoli Application Dependency Discovery Manager-discovered CIs, which contain detailed information about the items, including their relationships and dependencies. Tivoli Application Dependency Discovery Manager sensors *discover* (collect detailed configuration information from) the IT infrastructure, including identification of deployed software components, physical servers, network devices, virtual LAN, and host data used in a runtime environment.

This data includes CIs and relationships discovered using sensors, in addition to those alternately loaded through DLAs. This process is demonstrated in Figure 12-6.

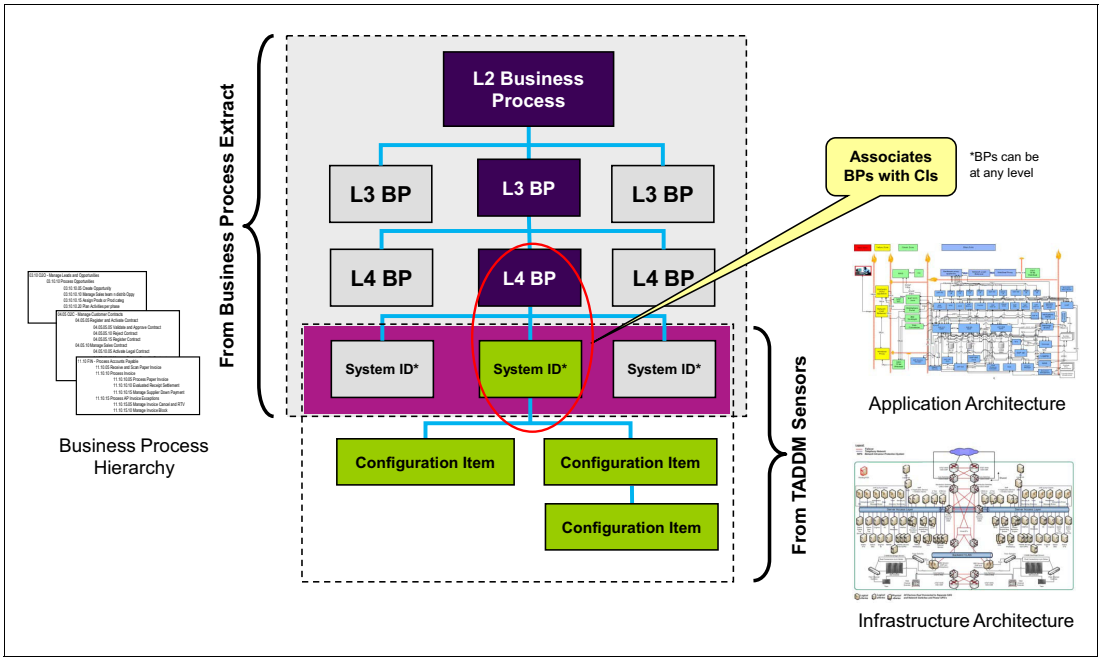

Figure 12-6 How the Business Process DLA works

In Figure 12-6, the enhanced BPH has been provided to the Business Process DLA. Figure 12-6 shows that a given Level 2 BP is composed of three Level 3 BPs. One of those Level 3 BPs, itself, consists of three Level 4 BPs. One of those Level 4 BPs has a dependency upon three different systems, as represented by their SIDs. All of that information was provided by the Enhanced BPH.

Meanwhile, Tivoli Application Dependency Discovery Manager already knows about the various SIDs and, in fact, is aware that one of the SIDs has two child CIs. One of these CIs has its own child CI, or technical subcomponent. The business process subject matter expert (SME) or application architect has no idea of the technical complexity of the system under the SID level.

However, by virtue of the BP DLA and an active Tivoli Application Dependency Discovery Manager implementation (in which all systems and applications in the enterprise are scanned), the business process SME has now effectively associated a Level 4 BP with an IT system and two levels of components underneath of it.

Depending upon the configuration of process criticality, and technical inheritance, which can be declared along with the enhanced BPH, a technical failure of that lower-level infrastructure component will not only flag its parent CI and SID, it will also signal the unavailability of that

Level 4 BP. Again, depending upon configuration, that Level 4 process availability might propagate all the way up to a Level 2 BP.

A Level 4 BP corresponds to an individual SAP transaction, or perhaps an interface. If a critical required system is down (for example, a messaging channel in the case of an interface), and that interface is defined as critical, the low-level unavailability of the messaging channel will cause that Level 4 and its parent BPs to be lit up in the service management tree visible in Tivoli Business Service Manager for the BPAM solution.

More information: The Business Process DLA is categorized as a utility, and is made available upon request to Tivoli Application Dependency Discovery Manager customers. To receive information about the source code and detailed instructions for the Business Process DLA, contact Derek Jennings at dmjennin@us.ibm.com or Mathew Davis at mdavis5@us.ibm.com.

The process of populating a Tivoli Application Dependency Discovery Manager database with a correlated set of BPs and SIDs described earlier is a useful technique. Modeling the BP and SID relationships directly in Tivoli Application Dependency Discovery Manager is possible, through its user interface (UI). However, if there are hundreds, or thousands of business processes, this task will be difficult.

Also, it would require that a Tivoli Application Dependency Discovery Manager SME have extensive business process knowledge, or a business process SME be adept at navigating and configuring Tivoli Application Dependency Discovery Manager.

12.5 Summary

IBM software portfolio includes several products with deep capability for providing systems management of large, complex business systems, as typified by SAP solutions. Any organization attempting to manage a large, SAP-centric enterprise system must take several factors into account when choosing their tools. In addition to examining how a product or product suite fits for a particular requirement, support organizations must also consider the comprehensiveness of coverage of the tools, ease of integration, and extensibility.

Heterogeneous systems require multiple levels of support and tools. However, to maximize availability, health, and stability, those tools should be able to be combined to provide an end-to-end, contextual view to support teams. This integrated worldview is increasingly important to business users and executives, who have little patience for arcane systems metrics that do not have correlation to their user experiences.

A business services management model is necessary to effectively marry a business process-centric and IT-centric view of system health, availability and performance. IBM enables a BSM model through BPAM, a solution that combines multiple IBM products. BPAM integrates these products in such a way as to minimize the need for manual configuration, while reducing the need for key personnel to take on new roles (business process SME to monitoring architect, or vice versa) just to enable a state-of-the-art solution.

12.6 References

These websites are also relevant as further information sources:

- IBM Tivoli Monitoring

  http://www.ibm.com/software/products/en/tivoli-monitoring-composite-app-mgmt

- IBM Tivoli NetView for z/OS

  http://www.ibm.com/software/products/en/tivoli-netview-zos

- IBM Tivoli Netcool/Impact

  http://www.ibm.com/software/products/en/tivonetc

- IBM Tivoli Netcool/OMNIbus

  http://www.ibm.com/software/products/en/ibmtivolinetcoolomnibus

- IBM Application Performance Management

  http://www.ibm.com/software/products/en/category/application-performance-management

- IBM Tivoli OMEGAMON XE for z/OS

  http://www.ibm.com/software/products/en/omegamon-xe-zos

- IBM SmartCloud Monitoring

  http://www.ibm.com/software/products/en/ibmsmarmoni

- Tivoli Network Manager IP Edition

  http://www.ibm.com/software/products/en/ibmtivolinetworkmanageripedition

- IBM Tivoli Business Service Manager

  http://www.ibm.com/software/products/en/tivoli-business-service-manager

- IBM Systems Director

  http://www.ibm.com/systems/director/

- IBM Tivoli Data Warehouse

  https://www.ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/Tivoli%20Documentation%20Central/page/Tivoli%20Data%20Warehouse

- IBM Tivoli Application Dependency Discovery Manager

  http://www.ibm.com/software/products/en/tivoliapplicationdependencydiscoverymanager

Related publications

The publications listed in this section are considered particularly suitable for a more detailed description of the topics covered in this book.

IBM Redbooks

The following IBM Redbooks publications provide additional information about the topic in this document. Note that some publications referenced in this list might be available in softcopy only.

- *Customizing and Extending IBM Content Navigator*, SG24-8055
- *Implementing Imaging Solutions with IBM Production Imaging Edition and IBM Datacap Taskmaster Capture*, SG24-7969
- *Integrating IBM Security and SAP Solutions*, SG24-8015
- *Introducing the IBM Security Framework and IBM Security Blueprint to Realize Business-Driven Security*, REDP-4528
- *Smarter Business: Dynamic Information with IBM InfoSphere Data Replication CDC*, SG24-7941

You can search for, view, download or order these documents and other Redbooks, Redpapers, Web Docs, draft and additional materials, at the following website:

ibm.com/redbooks

Other publications

This publication is also relevant as a further information source:

- *Enterprise Master Data Management: An SOA Approach to Managing Core Information, IBM Press,* ISBN-10: 0-13-236625-8; ISBN-13: 978-0-13-2366250

© Copyright IBM Corp. 2014, 2015. All rights reserved.

Online resources

These websites are also relevant as further information sources:

- IBM API Management

  http://www.ibm.com/software/products/en/api-management

- IBM Integration Bus

  http://www.ibm.com/software/products/en/ibm-integration-bus

- IBM Integration Bus Information Version 9.0

  http://www.ibm.com/support/knowledgecenter/SSMKHH_9.0.0/mapfiles/help_home_msgbroker.html?lang=en

- IBM InfoSphere Information Server Family

  http://www.ibm.com/software/data/integration/info_server/

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

Redbooks

IBM Software for SAP Solutions

SG24-8230-01

ISBN 073844104X

IBM

(0.5” spine)
0.475”<->0.873”
250 <-> 459 pages

SG24-8230-01

ISBN 073844104X

Printed in U.S.A.

Get connected