SAP Exchange Infrastructure 2.0

Performance Test jointly conducted by SAP AG and Sun Microsystems

Chapter 1

Executive Summary

As the number of IT systems supporting business processes grows, one-to-one communication links can become inefficient for integrating the systems to a coherent platform. By providing a platform for the integration of business processes, SAP Exchange Infrastructure 2.0 software enables companies to continue increasing the number of fully integrated business processes efficiently.

To provide customers with guidelines for deployment of SAP Exchange Infrastructure, SAP and Sun have teamed up in a performance test project. The project based measurements on a real-world integration scenario around mySAP Supplier Relationship Management (SRM) software. It also built on Sun's expertise in high-end system configurations, as well as in optimization of Java™ technology-based application performance.

The overall result of the performance tests was that SAP Exchange Infrastructure scales linearly with the hardware it was deployed on.

SAP Exchange Infrastructure 2.0 is available in ramp-up now, and unrestricted shipment is planned for mid-2003. The information provided in this document can help customers start planning their deployments now.

Chapter 2

Introduction

Today's core business processes are typically executed with IT support. Initially the IT department focused only on processes in central functions such as corporate financials or human resources (HR). But today, IT support is implemented for a multitude of functions including supplier relationship management (SRM), supply chain management (SCM), customer relationship management (CRM), and product lifecycle management (PLM). Often IT support in these areas is on separate systems, creating the challenge of integration to a single, consistent platform for execution of all core business processes an enterprise runs on.

With a small number of systems, set-up of one-to-one communication links is the easiest and fastest way to create that integrated platform. But as the number of systems grows, one-to-one links begin to create complexity in the landscape that could have a negative impact on operational cost and the stability of the entire IT environment.

Through a communication infrastructure such as the SAP Exchange Infrastructure—which allows central management of integration information combined with distributed execution—this impact can be avoided while the entire platform gains more flexibility to quickly accommodate changes in business requirements.

It is clear that resource requirements for SAP Exchange Infrastructure grow with the number of systems that connect, the number of messages sent between communication partners, and the size of the messages. After customers have gained initial experience with version 1.0 of SAP Exchange Infrastructure, they are now looking for a wider scope in deployment of SAP Exchange Infrastructure 2.0, requiring a guideline on scalability and resource requirements of the product under growing load.

In support of those needs, as well as a verification of the capabilities of SAP Exchange Infrastructure on the UNIX™ operating system, SAP and Sun have teamed up in a project to provide know-how on the operation of SAP Exchange Infrastructure in the high-end environments Sun is known for. But also, as SAP Exchange Infrastructure contains some major components developed using Java technology, Sun's expertise in optimizing Java application performance has been a key success factor of the project.

Chapter 3

General & Technical Descriptions: SAP Exchange Infrastructure and mySAP Supplier Relationship Management

The following chapters give a short general and technical overview of SAP Exchange Infrastructure, a component of SAP NetWeaver software — SAP's comprehensive integration and application platform — and the mySAP Supplier Relationship Management (SRM) solution. This will help prepare the reader for the more specific descriptions of the performance test.

3.1 SAP Exchange Infrastructure: General & Technical Description

SAP Exchange Infrastructure is the layer of SAP NetWeaver for integrating heterogeneous software components of a system landscape. For this purpose, SAP offers an Integration Server and different adapters to connect multiple business systems through SAP Exchange Infrastructure.

3.1.1 Overview

The collaborative nature of e-business mandates a decentralized and distributed structure to provide the necessary flexibility. From a company's point of view, the distributed component-based landscape should not be visible at all to a user. The company's business processes constitute the technical linkage of the components. With the help of SAP Exchange Infrastructure it is possible to implement cross-system business processes by integrating different versions of both SAP and non-SAP systems implemented on different platforms (Java, ABAP, etc.).

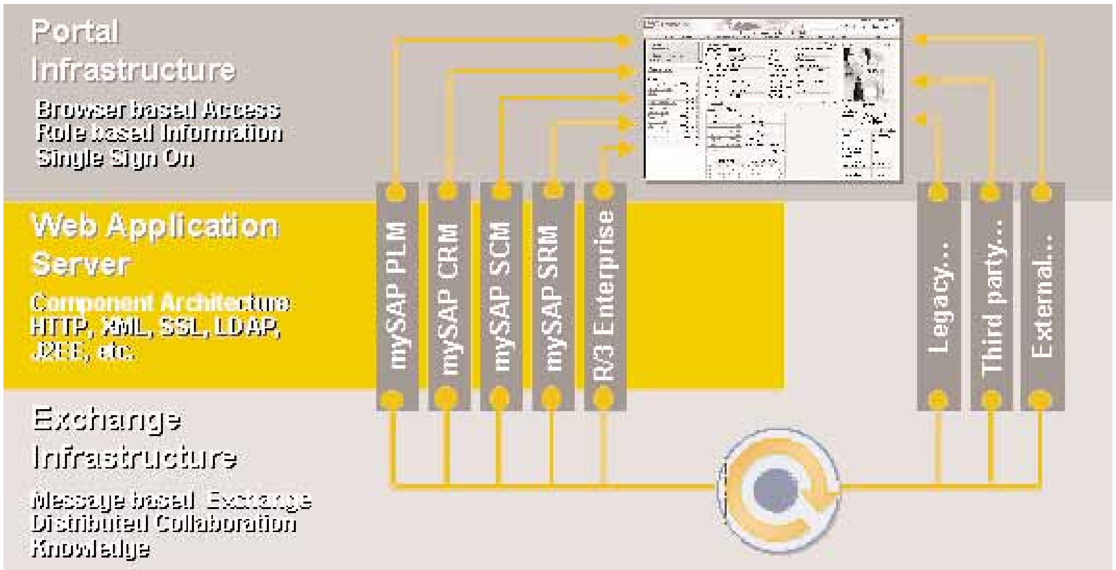

Figure 1: Building blocks

SAP Exchange Infrastructure is based on an open architecture and makes use of open standards such as extensible Markup Language (XML), Web Service Description Language (WSDL), and Simple Object Access Protocol (SOAP). It also offers services that are essential in a heterogeneous and complex system landscape—namely a runtime infrastructure for message exchange, configuration options for managing business processes and message flow, and options for transforming message contents between the sender and receiver systems (mapping).

3.1.2 Architecture of SAP Exchange Infrastructure

The SAP Exchange Infrastructure product has the following core building blocks:

- Integration Repository to capture integration-relevant knowledge that can be determined at design time.
- Integration Directory, which drives the collaboration execution by containing the knowledge to describe integration-related parts of the configured customer landscape.
- Integration Server, which contains the centrally configured Integration Engine as well as further integration services, and is only used for integration purposes. It receives XML messages, determines the receiver, performs mappings, and routes the XML messages to the respective receiver systems. In doing so, it is dependent on integration information that is stored in the Integration Directory.
- Integration Monitor, which helps deal with technical and business-driven monitoring of the exchange infrastructure.

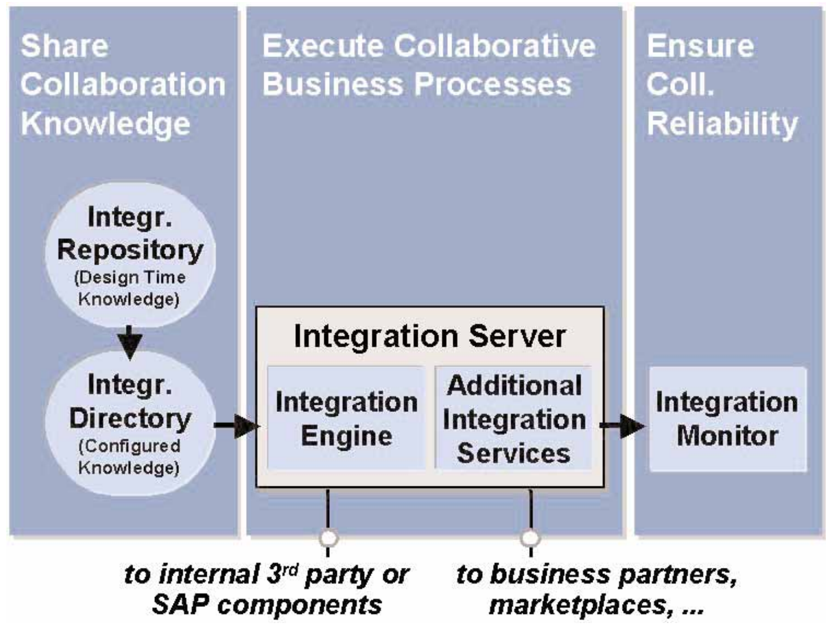

Figure 2: Overview of SAP Exchange Infrastructure

3.1.3 Message Flow within the Integration Engine

The execution of processes takes place in the runtime of SAP Exchange Infrastructure. The Integration Server and particularly the Integration Engine are the parts of SAP Exchange Infrastructure that cause the system load.

The Integration Engine is a central “distribution engine” that processes XML messages, regardless of whether a message was sent to the Integration Engine using an adapter or an application proxy. This includes services for determining receivers (logical and technical routing) and for the transformation of message contents between sender and receiver systems (mapping). Without those components it would not be possible to exchange messages using SAP Exchange Infrastructure.

Figure 3 illustrates a single message being transferred through the Integration Engine between two applications. Routing and physical address resolution is only needed for the request as the response is transferred to a sender, which is already known.

Different kinds of adapters ensure connectivity to business partners, third-party systems and SAP solutions.

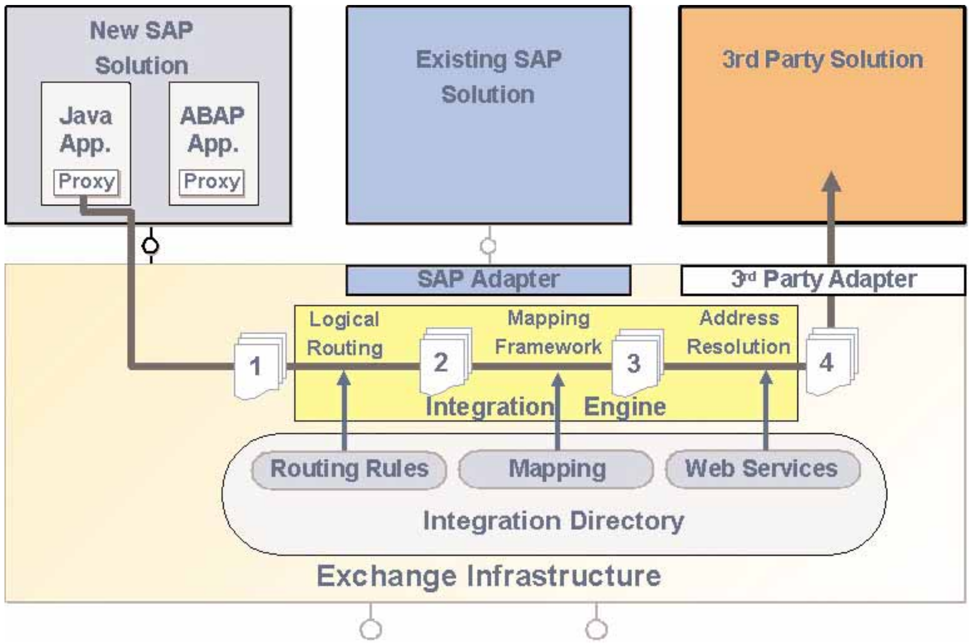

Figure 3: Overview of the process flow within the Integration Engine

Given a message, including information about the sender and the message interface, the **logical routing** service determines receivers and required interfaces by evaluating the corresponding routing rules, whether those rules are XPath expressions or Java code. This logical routing can have a significant influence on the overall performance of the Integration Engine.

The **mapping service** determines the required transformations that depend on message, sender, and sender interface, as well as the receiver and receiver interface. In the case of synchronous communication, even the message direction is relevant to appropriately transform input, output, and fault messages. After retrieving the required mapping from the Integration Directory, the mapping service can then either execute XSLT mappings or Java code (or any combination in a given sequence) to the business content of the sender message. Mapping, like logical routing, signifies changes to the data structure and therefore can have an impact on performance.

3.2 mySAP Supplier Relationship Management: General & Technical Description

Supporting your e-procurement and sourcing activities with mySAP Supplier Relationship Management (SRM) software helps to cut the costs of goods sold and rationalize the supply base, and it can deliver quick return on investment (ROI). It can also help drive process extension and supplier collaboration by automating processes with all suppliers for all purchased goods and services company-wide.

The complete mySAP SRM solution consists of the following key focus areas and capabilities:

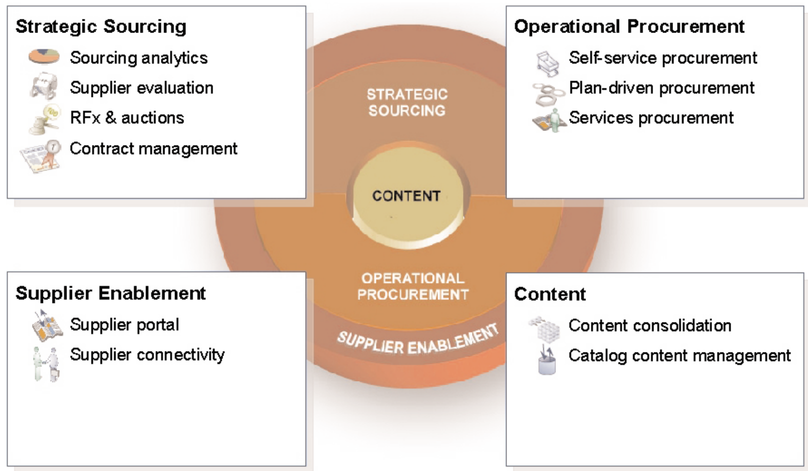

Figure 4: mySAP SRM components, features, and capabilities

mySAP SRM is the first SAP solution that is entirely built on top of SAP's new application and integration platform, SAP NetWeaver. SAP Exchange Infrastructure, a component of SAP NetWeaver, is used by the various application components of mySAP SRM as a common infrastructure for process integration. Figure 5 below shows all mySAP SRM application components and how they integrate with SAP Exchange Infrastructure as well as with the back-end systems of the buying company.

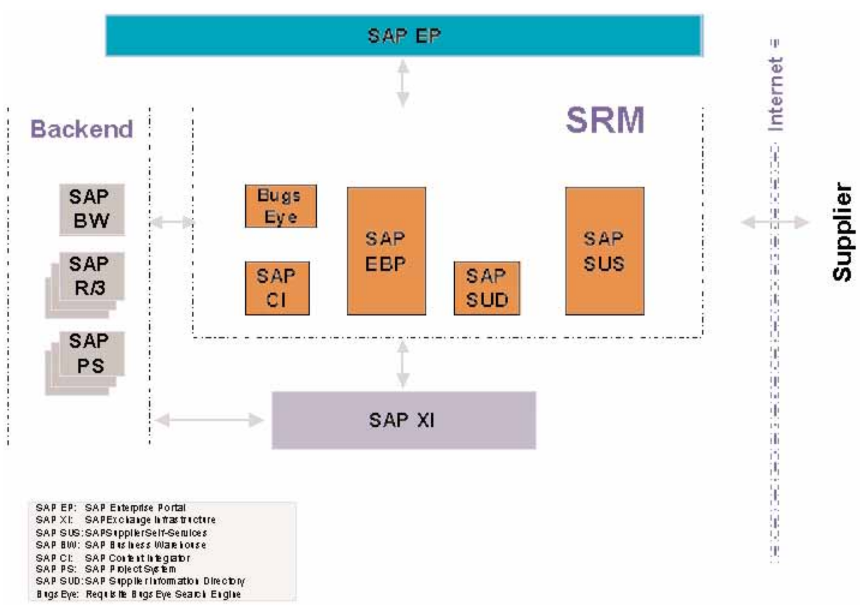

Figure 5: SRM application components and integration with SAP Exchange Infrastructure and SAP Enterprise Portal

For the performance test three business processes were chosen. Two of them, Purchase Order (PO) exchange and Trading Partner (TP) replication, belong to key focus area "Supplier Enablement" and the third process called Catalog Import belongs to the Catalog Content

Management capabilities of the key focus area “Content.” This section gives a more detailed overview of these key capabilities and processes.

3.2.1 Supplier Enablement

The Supplier Enablement capabilities of mySAP SRM software allow purchasing organizations to set up a Web-based front end with self-services for suppliers. The services provide suppliers with a streamlined order management system, along with a user-friendly interface and comprehensive search options that help make the job easy. It allows buyers to benefit from the advantages of e-procurement. Supported processes encompass creating and updating catalog content, processing of orders and invoices by the supplier, and the collection of data about performed services. Supplier Enablement capabilities are delivered through SAP Supplier Self-Services (SAP SUS), one of the application components of mySAP SRM software.

The supplier enablement scenario normally looks like this:

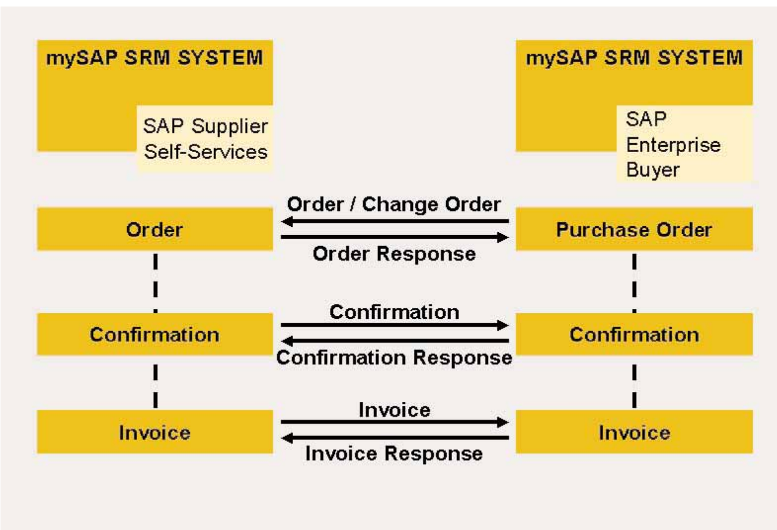

Figure 6: Typical Supplier Enablement scenario

The first step in the shown business scenario is the Purchase Order (PO) exchange process used in the performance test. Here a purchase order message is sent from SAP Enterprise Buyer, the core e-procurement application of mySAP SRM, to SAP Supplier Self-Services via SAP Exchange Infrastructure.

The second business process chosen for this performance test, **Trading Partner (TP) replication**, also belongs to this key focus area. The relevant business data about a new supplier is again exchanged between SAP Enterprise Buyer and SAP Supplier Self-Services. The interesting aspect of this process is that the message can be routed directly from the sending to the receiving interface without any mapping step required.

3.2.2 Catalog Content Management

Catalog Content Management is the process of managing the flow of content that supports procurement activities. Usually, content revolves around suppliers and product data (often in the form of a searchable catalog).

Chapter 4

Business Processes Used in the Test Environment

The picture below shows the software architecture that composed the performance test environment. These software components and in particular the required components of the mySAP SRM solution were chosen as they constitute the first predefined SAP Exchange Infrastructure scenarios available for customers.

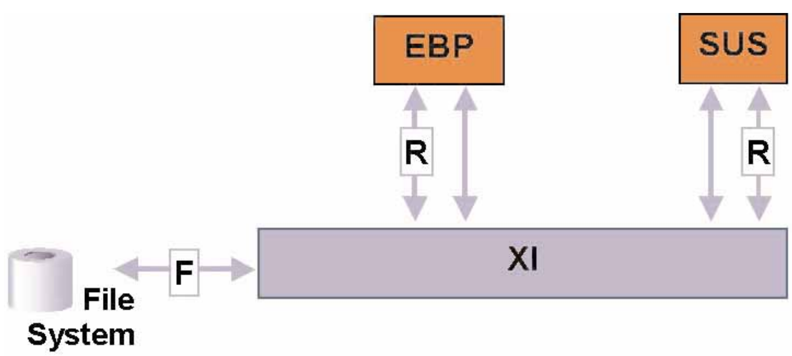

Figure 7: Performance test environment software architecture

- *SAP Enterprise Buyer:* SAP Enterprise Buyer is the core e-procurement component of mySAP SRM. As such, it enables employees, managers, and professional buyers to execute their individual e-procurement tasks using a Web-based interface tailored to their individual roles. For example, employees can use self-service procurement functionality for centralized direct procurement, and professional purchasers can leverage SAP Enterprise Buyer to optimize purchasing decisions. SAP Enterprise Buyer is an ABAP-based application running on SAP WebApplication Server (SAP WebAS).
- *SAP Exchange Infrastructure:* SAP Exchange Infrastructure is used for document exchange between application components of mySAP SRM, between mySAP SRM components and back-end systems of the purchasing organization and connectivity to supplier systems, respectively. In general, SAP Exchange Infrastructure drives process-centric integration and collaboration

by integrating business processes and systems. SAP Exchange Infrastructure provides a message-oriented infrastructure based on open Internet standards, such as XML and Simple Object Access Protocol (SOAP). In order to integrate existing applications and even legacy systems SAP Exchange Infrastructure supports conventional communication protocols via dedicated adapters, e.g. Remote Function Calls (RFCs, symbolized by the "R" in the figure) via RFC Adapter.

The file system represents the storage that holds the new or updated catalog content as part of the Catalog Import process. Access to the file system is done via a File Adapter (symbolized by the "F" in the figure).

4.1 Measured Business Processes

This chapter describes the three business processes investigated in this study.

4.1.1 Process: Purchase Order Exchange

This business process is part of all mySAP SRM procurement scenarios. The relevant steps of the entire scenario, which were used in the performance tests, are 1) a local purchase order is created for goods or services in SAP Enterprise Buyer (EBP), the e-procurement system of mySAP SRM; 2) the purchase order is sent via SAP Exchange Infrastructure to SAP Supplier Self-Services (SUS) component, triggering the creation of a sales order. In addition it is possible to create confirmations and invoices from orders. The figure below shows this process and the components used.

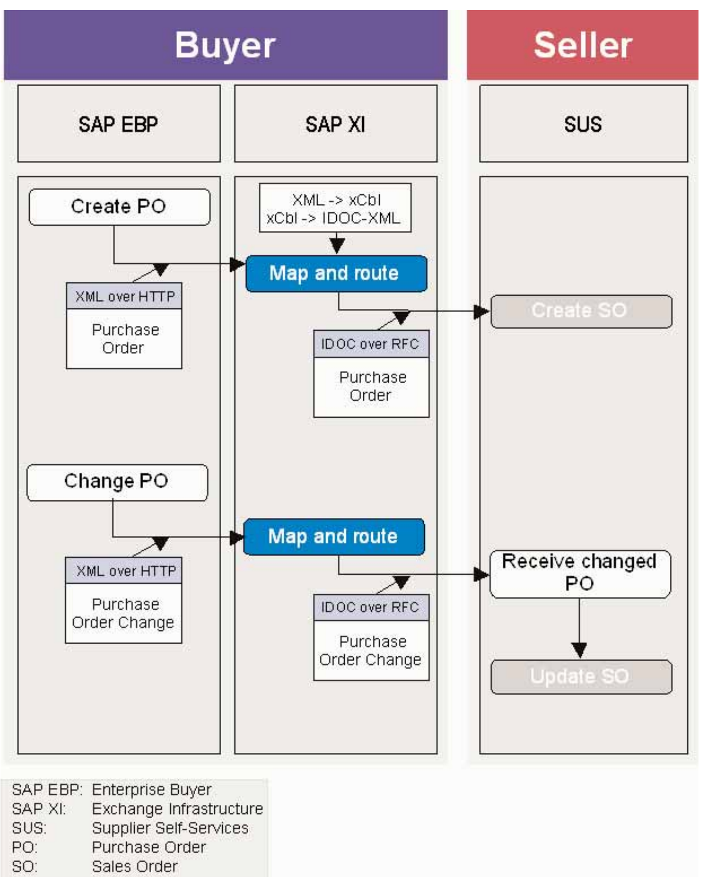

Figure 8: The purchase order exchange process

4.1.2 Process: Trading Partner Replication

Within mySAP SRM the central place to store and maintain Trading Partner information is SAP Enterprise Buyer (EBP). A Trading Partner (more commonly known as a business partner object in SAP Enterprise Buyer) holds all relevant data of a trading partner. This information is shared with the other mySAP SRM application components, in particular with SAP Supplier Self-Services (SUS). Trading Partner replication is the business process that replicates trading partner information between SAP Enterprise Buyer and SAP Supplier Self-Services via SAP Exchange Infrastructure. The figure below shows this process and the components used.

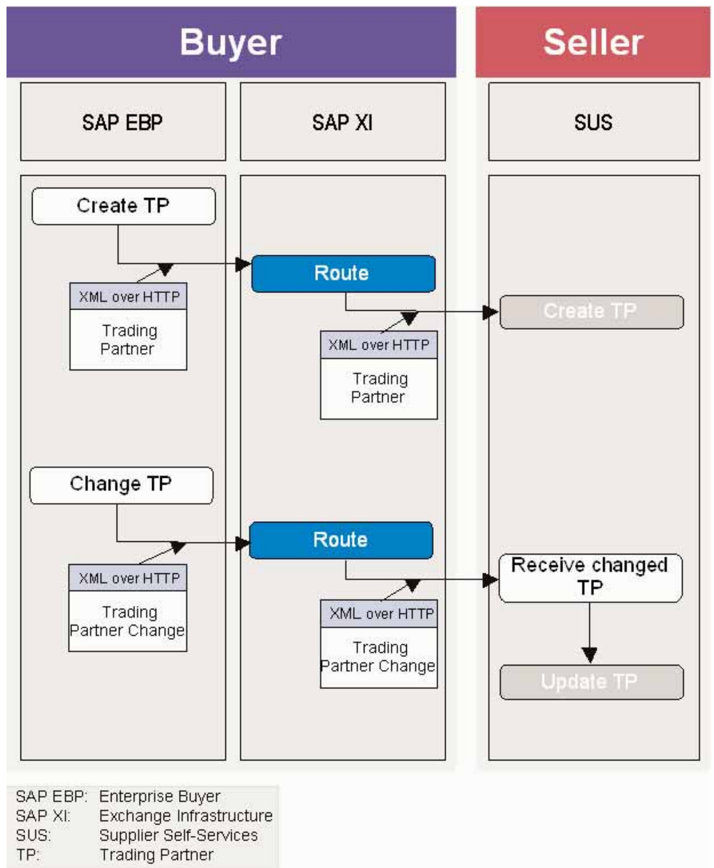

Figure 9: Trading partner replication process

4.2 Business Scenario: Catalog Content Management

Using SAP Content Workbench, a module of SAP Enterprise Buyer providing an XML-based interface for uploading product content from external business partners, system administrators can import external catalogs and map the external category and attribute structures to the internal structures. Before the data can be passed to the content workbench, the mapping of the imported product data must be completed.

4.2.1 Process: Catalog Import

External product catalog data or product catalogs are uploaded into the SAP Enterprise Buyer system using an XML-based interface. The incoming XML documents are processed using SAP Exchange Infrastructure. Supported XML schemata are mapped to the SAP EBP schema. The following open and industry-specific standards are supported: BMEcat 1.01, xCBL 3.0, eCx XML 2.0. After mapping, the data is transferred to the SAP Enterprise Buyer system:

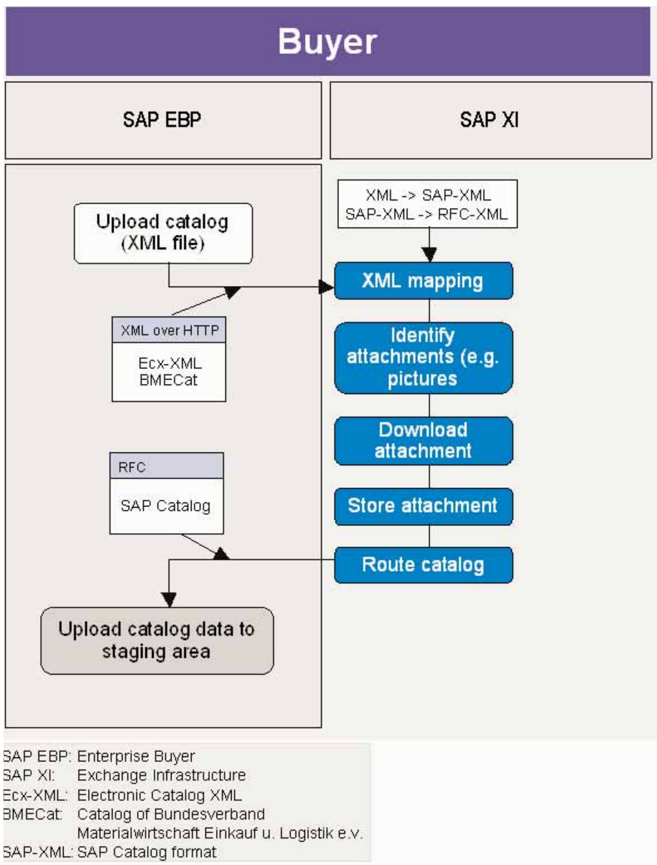

Figure 10: Catalog import process

Chapter 5

Technical View

The SAP Exchange Infrastructure performance test was conducted at the Sun Microsystems Benchmark Center in Langen, Germany, on a Sun Fire™ 15K server.

Sun Microsystems has a long history of high performance and scalability achievements with SAP's Standard Application Benchmarks and industry-specific performance tests. The SAP Exchange Infrastructure performance test gave SAP and Sun the opportunity to expand on that experience to include SAP Exchange Infrastructure, where in-depth knowledge of Java technology is essential, since some parts of SAP Exchange Infrastructure are Java technology-based.

5.1 Hardware setup

The entire performance test was run on one Sun Fire 15K server. Thus, the actual SAP Exchange Infrastructure, the SAP EBP-SAP SUS system acting as message consumer, and the file adapter used as benchmark driver were co-located on the same server.

Server configuration:

Model	Sun Fire 15K
CPUs	72 UltraSPARC™ III, 1050-MHz, 8 MB L2 cache
Memory	288 GB
Storage	9 Sun StorEdge™ T3 arrays
SAPS[1]	29820 [2]

A total of nine Sun StorEdge T3 storage arrays were used for the performance test. These units are equipped with nine internal disks each. Thus, a total of nine separate FCAL I/O channels and 81 disks were available, allowing an efficient distribution of the I/O load. In addition the storage units are equipped with a large write cache to speed up I/O operations.

---

1 SAPS is a throughput unit that describes the power of a configuration. Details can be found on the SAP Service Marketplace at http://www.sap.com/benchmark. 100 SAPS are defined as 2,000 fully business processed order line items per hour in the standard Sales & Distribution (SD) application benchmark. This throughput is achieved by processing 6,000 dialog steps (screen changes) and 2,000 postings per hour in the SD benchmark or processing 2,400 SAP transactions.

2 SAPS value achieved with 72-CPU Sun Fire 15K server containing 1200-MHz CPUs.

Since the main load was related to the SAP Exchange Infrastructure system's database, wide thin stripes [Lar2000] spanning six of the T3 units were used for the main data volumes of this database. In addition, the redo logs were placed on a separate T3 unit to ensure a dedicated I/O channel for these logs.

Detailed storage subsystem configuration:

Component	T3/1	T3/2	T3/3	T3/4	T3/5	T3/6	T3/7	T3/8	T3/9
SAP Exchange Infrastructure Database: data, sys	x	x	x	x	x	x			
SAP Exchange Infrastructure Database: log							x		
SAP EBP Database: data, sys, log								x	
SAP Exchange Infrastructure, SAP EBP: /usr/sap, /sapmnt									x
Backup									x

Figure 11: SunFire 15K server with attached Sun StorEdge T3 array storage

5.2 Software setup

Software component	Version
Operating environment	Solaris™ OE 9 Rel. 9/02
SAP Exchange Infrastructure	2.0
SAP Exchange Infrastructure Database	SAP DB 7.3.0.29
mySAP SRM	2.0
SAP Enterprise Buyer	3.5
Database for SAP Enterprise Buyer	SAP DB 7.3.0.29
Sun HotSpot™ Java Virtual Machine	1.3.1_06

The SAP Exchange Infrastructure performance test demonstrated and confirmed once again the capability of the Solaris Operating Environment to efficiently utilize systems with very large numbers of CPUs. Well-established features such as processor sets, Solaris Resource Manager (SRM) software, and the use of domains allow for effective resource allocation and control for several co-located applications components on very large systems. This performance test also demonstrated that new Solaris capabilities like Memory Placement Optimization (MPO) [MPO2002] further increase the performance and scalability especially of such large systems.

SAP DB is SAP's reliable, high-performance database, developed to meet the demands of high-volume online transaction processing (OLTP) that is typical for e-business and applications such as SAP Exchange Infrastructure. SAP DB provides a multi-process/multi-threaded server architecture, SMP scalability, and minimal I/O operations.

5.3 Performance Test Procedure

Since no official SAP Standard Application Benchmark has been established yet for SAP Exchange Infrastructure, a performance test procedure was defined within the scope of this project. This procedure was designed to be closely aligned with the standard SAP benchmark procedure used for example in the well-known SAP Sales & Distribution (SAP SD) benchmark:

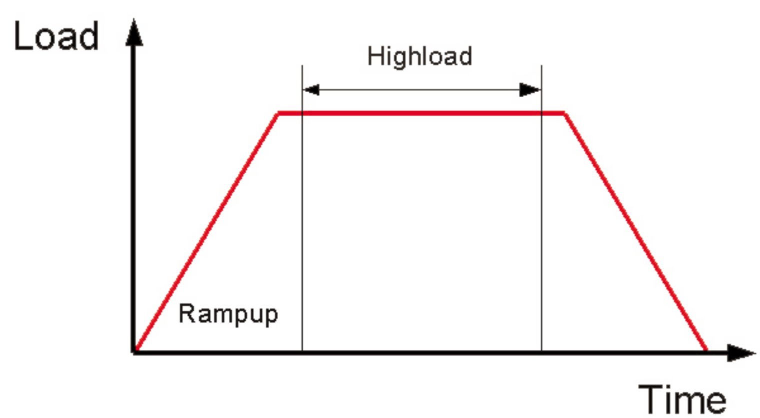

Figure 12: Load profile of a performance test

- *Step 1:* The two SAP systems (SAP Exchange Infrastructure and EBP) are started up. In order to generate a high load on the SAP Exchange Infrastructure, the EBP system acting as the document sender as described in figures 8, 9, and 10 is replaced with a set of file adapters. These adapters simulate many EBP systems sending messages to SAP Exchange Infrastructure in parallel. On the consumer side the SAP EBP-SAP SUS receives the messages but does not process them. This is because of the fact that the focus is on SAP Exchange Infrastructure and the load of the receiving system should be as low as possible.
- *Step 2:* The file adapter performance drivers are started up sequentially. They start to generate load by sending XML documents to the inbound ICM (Internet Communication Manager of the SAP Web Application Server) processes of the SAP Exchange Infrastructure instances. A varying number of file adapters are used depending on the number of SAP Exchange Infrastructure processing instances to generate load—typically a single adapter produces enough load for one processing instance. Since the adapters are started sequentially, the load increases gradually during this ramp-up phase. In the environment developed for this project, the duration of the ramp-up phase is configurable, and a value of 120 seconds was found suitable.
- *Step 3:* After all file adapters are started up and operational, the load reaches a plateau where it remains fairly constant. Soon after this happens, the so-called high-load phase begins. A high-load duration of 900 seconds was chosen here, which is in sync with the requirements of SAP Standard Application Benchmarks. This period of time is a good compromise showing both stability for more than just a few minutes while being well suited for practical usage in terms of turnaround times, where run durations of several hours would not be acceptable.
- *Step 4:* After the 900 seconds of the high-load phase the file adapters are stopped and the system load decreases while messages already in the message queues are processed. Eventually, the performance test is completed once there are no more messages in the system. During the high-load phase, the length of the message queue was targeted to remain fairly constant to prove that the system has reached a stationary state.

ABAP reports developed during this project then collect the performance data for the messages that were processed and compute time-dependent as well as total throughput values for the high-load phase. In addition, numerous other Solaris and third-party tools (such as vmstat, iostat, mpstat or vxstat) were used to collect performance data on the system, processor set, and process level during the run. This data is crucial to understanding the system behavior and to identify areas for potential performance improvements. All SAP Java 2 Platform, Enterprise Edition (J2EE™) Engine servers were run with garbage collection tracing and procedures have been developed to analyse this output to derive both the performance impact this activity has, as well as recommendations how to configure the Java Virtual Machine to improve performance.

5.4 Tuning Steps

The scalability and performance of the SAP Web Application Server architecture, combined with Sun servers, has been demonstrated time and again in benchmarks as well as at thousands of customer sites. SAP Exchange Infrastructure adds a new dimension to this well-proven architecture by including J2EE applications to the process flow. Consequently, the tuning areas looked at for this project comprise well-known ones as well as some new ones:

- *Instance setup:* The number and configuration of SAP instances is an important factor for the scalability and overall performance of an application. The main parameters to consider are:

  Number of SAP instances

  Number of work processes (DIA) per instance

  Number of SAP J2EE Engine servers per instance

- *Java tuning:* The Sun HotSpot Java Virtual Machine (JVM) is a highly configurable environment capable of supporting a wide spectrum of different application types ranging from small, short-lived client applications to server applications running for days or weeks, requiring gigabytes of heap space. The main areas to look at in order to obtain the best performance for a given application load are the total heap size and the sizes of the different heap generations (young, old, permanent).
- *Database tuning* is a well-established procedure in performance tests and benchmarks as it is in daily operations. Typical areas to consider include buffer cache sizes (such as data cache, log buffer), savepoint performance and configuration changes leveraging SMP scalability.
- *I/O tuning* was limited to the use of wide thin stripes and dedicated I/O channels for crucial files as described in section 5.1.

The Solaris Operating Environment was designed for performance and scalability even for such large numbers of CPUs as were used in this performance test. New performance-related features have been added over time and more will be added in the future. Examples are the use of processor sets [Laux2001], process priority control, or the recently added Memory Placement Optimization (MPO), which directly improves process scheduling and memory management transparently to the application.

Chapter 6

Results

This chapter summarizes the performance results obtained for the different business processes investigated in this study.

6.1 Purchase Order Exchange Scenario

From a customer perspective this scenario is the most important one as purchase orders are deemed to represent the customer's daily usage of mySAP SRM. From a technical point of view the PO scenario includes two mappings inside the SAP Exchange Infrastructure engine, which creates a system load representative of real life behavior of SAP Exchange Infrastructure in customer environments.

Consequently, most of the project effort went into gathering data on performance and scalability for this type of load. Several different message sizes (in terms of line items per PO) were investigated up to messages containing 3000 line items, and system scalability using different numbers of CPUs and SAP instances was demonstrated.

6.1.1 Throughput results

Figure 13 shows the system throughput in units of line items per time interval. As one would expect, the throughput increases significantly as the number of line items per message increases since the processing and routing overhead required for each message is considerably higher than the processing time required for each individual line item:

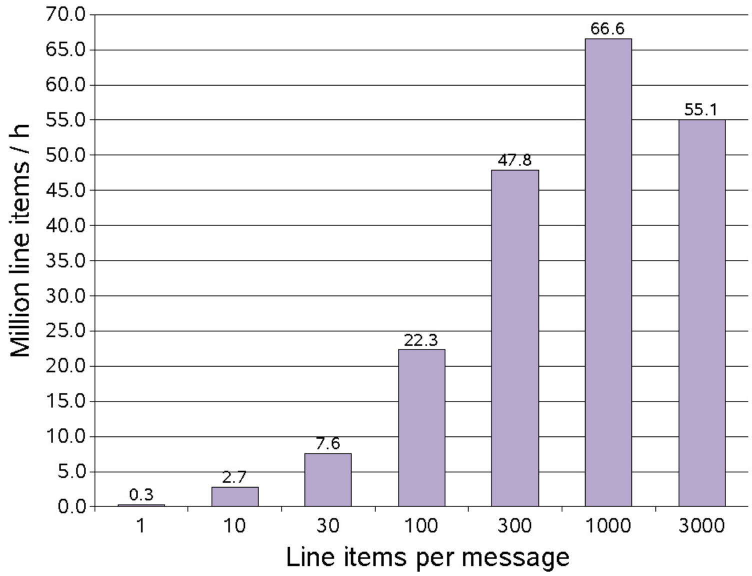

Figure 13: System throughput for different message sizes

In order to minimize configuration changes in production systems, one goal of this project was to find a system configuration suitable to efficiently handle different message sizes. The runs shown in figure 13 used the same configuration except for the young generation size of the Java Virtual Machine, which was set at 256 MB for the smaller message sizes (300 line items and below) and at 512 MB for the larger message sizes in order to reflect possible scenarios in customer environments. This accounts for the different requirements imposed on the Java environment by the varying number of line items, and is reflected in the Java garbage collection times which increase from below 2% for the smaller messages up to well above 10% for the larger messages. Again, using a specific system configuration with varying Java generation sizes and different numbers of worker threads for each message size would certainly reduce these times for larger messages, but this in turn would be impractical in real-life applications of such an infrastructure.

The following configuration was used for all these runs:

Component	CPU#
SAP Exchange Infrastructure Processing Instances[1]	63
SAP Exchange Infrastructure Database	5
File Adapter	2
EBP	2

1 The term “processing instance” is used here to describe what is typically called “application instance” in SAP 3-tier client-server environments to distinguish it from the database instance. Since SAP Exchange Infrastructure is not a business application, the term “application instance” could be confusing.

6.1.2 Scalability results

Figure 14 shows the excellent scalability results for the purchase order scenario, using two different message sizes as examples. In each case, the number of SAP instances was varied between 1, 2, 4, 8, and 16, and the number of CPUs used for the SAP Exchange Infrastructure processing instances was constrained accordingly.

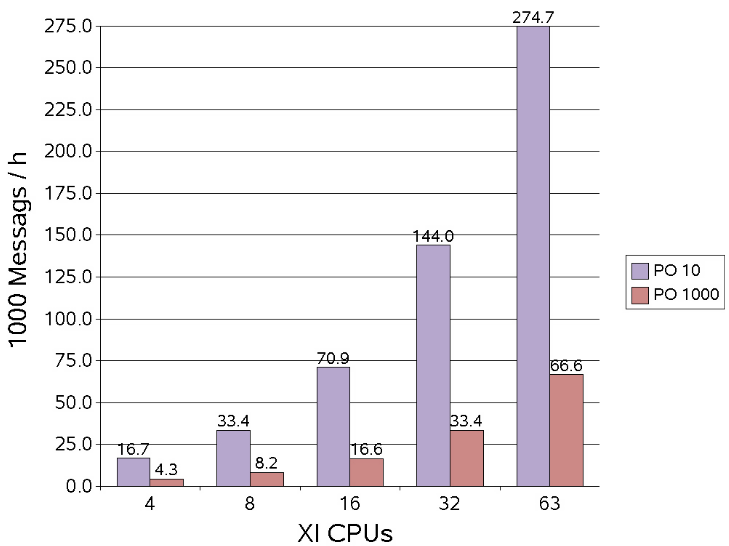

Figure 14: System scalability for PO processing scenario with different line item numbers

6.2 Trading Partner Scenario

This particular scenario does not involve a mapping, so document routing remains the main task of the SAP Exchange Infrastructure.

6.2.1 Throughput results

The maximum throughput achieved for this scenario was more than 440,000 messages per hour. The set-up for this result used 15 SAP Exchange Infrastructure processing instances and the following configuration:

Component	**CPU#**
SAP Exchange Infrastructure Processing instances	46
SAP Exchange Infrastructure Database	9
File Adapter	3
EBP	14

Note that the load generated by the SAP EBP system was considerably higher in this scenario; thus the available CPU resources for the SAP Exchange Infrastructure system were less than in the purchase order scenario.

6.2.2 Scalability results

Figure 15 below shows the scalability behavior of the system using 1, 2, 4, 8, or 12 SAP Exchange Infrastructure processing instances. The size of a common trading partner message (changing only one trading partner at a time) is 1 Kb. In order to analyze a higher system load a second message with 300 Kb of size was used in addition:

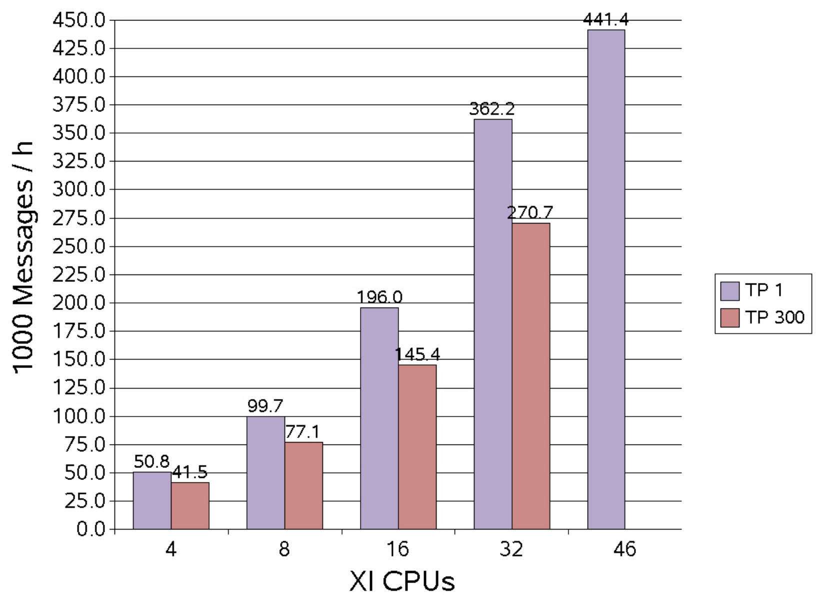

Figure 15: System scalability for TP scenario with different line item numbers

Again, the excellent scalability of the system is demonstrated. Note also that the influence of the message size is smaller here than compared to the purchase order scenario, which corresponds to the absence of any mapping requiring CPU power in this scenario.

6.3 Catalog Scenarios

During the performance test two different catalogs were used:

1. BMECat format with approximately 7 MB XML data
2. eCX format with approximately 30 MB XML data

These catalog scenarios have different characteristics than the purchase order (PO) and the trading partner (TP) scenario due to their message size and the increased processing requirements per message. While PO and TP messages typically take well below 1 second to be processed, these catalogs required processing times in the order of several minutes. It was observed that this processing time is spent almost 100% in the Java Virtual Machine without any noticeable garbage collection or system activity. In addition, there was only a small amount of I/O activity, so almost the entire system could be used for SAP Exchange Infrastructure processing instances:

Component	CPU#
SAP Exchange Infrastructure Processing instances	69
SAP Exchange Infrastructure Database	1
File Adapter	1
EBP	1

The throughput results are summarized in this table:

Catalog	Size [MB]	Proc. Time [s]	Throughput [Cat/h]	Throughput [MB/h]
BMECat	7	632	392	2840
eCX	30	223	1065	32280

As can be readily seen from the processing times, the resource requirements for the two catalog formats are very much different, as the eCX catalog was processed much faster than the BMECat catalog despite it being represented by more than four times the XML data in size. The difference in processing time for the different catalog formats is explained by the more complex mapping required to map the BMECat format to the SAP Enterprise Buyer catalog format.

6.4 Linear Scalability

In this project several scalability scenarios have been evaluated. The scenarios were run by varying the number of SAP Exchange Infrastructure processing instances (and constraining the number of CPUs per instance accordingly), and by varying the message size.

As an overall result the system shows an excellent and linear scalability. This scaling property is also valid for message processing, including complex mappings on the SAP J2EE Engine.

Chapter 7

Conclusions

The results of the performance test prove SAP Exchange Infrastructure 2.0 and the Sun Fire 15K server are capable of mapping and routing large throughputs of messages, resulting in the huge number of over 66 million mapped and routed line items per hour. The scalability results were impressive. The combination of SAP Exchange Infrastructure 2.0 software and the Sun Fire 15K hardware showed near linear scalability from 1 up to 63 CPUs used for the SAP Exchange Infrastructure processing (note that the other CPUs in the server were used for different tasks like SAP Exchange Infrastructure database and file adapter). Furthermore, the ability to scale has been demonstrated for the range of applications ranging from simple message routing to message processing including complex mappings on the SAP J2EE Engine.

This power makes SAP Exchange Infrastructure 2.0 and the Sun Fire 15K server the perfect platform to meet exchange infrastructure needs in high-performance environments where scalability is critical to solve current and future requirements.

The test included real-life scenarios as well as scenarios that were chosen to demonstrate the scalability of the tested exchange environment by increasing the number of messages as well as the number of line items per messages. The usage of the different message sizes was also important to assure that productive systems can be configured to efficiently handle different message sizes without any changes since reconfigurations and different tuning methodologies for different message sizes would be impractical in real-life usage of such an exchange infrastructure. The results of the performance test of the combination of SAP software and Sun hardware described in this document exceeded the goals established at the beginning of the test. During the test several key success factors for those results were identified:

Key success factors

- Multitier architecture of SAP applications
- Software scalability
- Hardware scalability
- Solaris Operating Environment
- Instance setup
- Java tuning
- Database and I/O Tuning

Multitier Architecture of SAP applications

With SAP's multitier architecture and the capabilities of the Sun Fire 15K server to support this architecture, the workload of the different application components (SAP Exchange Infrastructure processing instances, SAP Exchange Infrastructure Database and File Adapter) could easily be distributed among an optimum number of processors depending on the tasks that needed to be fulfilled.

Software Scalability

Deploying more SAP Exchange Infrastructure processing power by increasing the number of SAP Exchange Infrastructure processing instances shows the ability of SAP Exchange Infrastructure 2.0 to scale. In addition by deploying those instances on the same machine demonstrates the ability to scale on a large server.

Hardware Scalability

The Sun Fire 15K server used as the hardware platform during the performance test facilitates a highly flexible system configuration with all application components running on one physical server. This enabled online reconfiguration of the components on different numbers of CPUs, such as increasing the CPU usage of one SAP Exchange Infrastructure processing instance from four to eight CPUs without affecting the other application components and without reboot.

Solaris Operating Environment

The Solaris Operating Environment was designed for performance and scalability to a large number of CPUs as were used during the performance tests. It facilitates performance-related features such as processor sets, process priority control and the Memory Placement Option, which directly enhances process scheduling and memory management transparently to the application.

Instance Set-up

The number and configuration of SAP instances is crucial to the scalability and performance of the SAP application. Consequently it is important to consider the number of SAP instances with the number of dialog processes and the number of SAP J2EE Engine servers per instance.

Java Tuning

As the Sun HotSpot Java Virtual Machine (JVM) is a highly configurable environment used for a wide spectrum of different application types the proper Java tuning becomes crucial for the overall performance and scalability of the exchange infrastructure. To obtain the optimal performance for a given application load, the total heap size and the sizes of different heap generations have to be optimized.

Database and I/O Tuning

During performance tests as well as in daily productive operations the database tuning due to the individual application load is critical to obtain the optimal performance. The required tuning steps are database specific and need to account for the underlying I/O subsystem.

Chapter 8

References

[Lar2000] B. Larson: *Wide Thin Disk Striping.* Sun BluePrint.
URL: http://www.sun.com/solutions/blueprints/1000/layout.pdf

[Laux2001] M. Laux: *Solaris™ Processor Sets Made Easy.*
URL: http://soldc.sun.com/articles/solaris_processor.html

[MPO2002] *What's New in the Solaris 9 Update Releases?* Sun Microsystems.
URL: http://www.sun.com/db/doc/816-7173/6md6rliq8?q=memory+placement&a=view

[Sun2003] Performance Documentation for the Java HotSpot Virtual Machine. Sun Microsystems.
URL: http://java.sun.com/docs/hotspot/

Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 USA 1-650-960-1300 or 1-800-555-9sun www.sun.com

We make the net work.

Sun Worldwide Sales Offices: Argentina: +5411-4317-5600, Australia: +61-2-9844-5000, Austria: +43-1-60563-0, Belgium: +32-2-704-8000, Brazil: +55-11-5187-2100, Canada: +905-477-6745, Chile: +56-2-3724500, Colombia: +571-629-2323, Commonwealth Of Independent States: +7-502-935-8411, Czech Republic: +420-2-3300-9311, Denmark: +45 4556 5000, Egypt: +202-570-9442, Estonia: +372-6-308-900, Finland: +358-9-525-561, France: +33 134-03-00-00, Germany: +49-89-46008-0, Greece: +30-1-618-8111, Hungary: +36-1-489-8900, Iceland: +354-563-3010, India: Bangalore: +91-80-2298989/2295454; New Delhi: +91-11-6106000; Mumbai: +91-22-697-8111, Ireland: +353-1-8055-666, Israel: +972-9-9710500, Italy: +39-02-641511, Japan: +81-3-5717-5000, Kazakhstan: +7-3272-466774, Korea: +822-2193-5114, Latvia: +371-750-3700, Lithuania: +370-729-8468, Luxembourg: +352-49 11 33, 1, Malaysia: +603-21161888, Mexico: +52-5-258-6100, The Netherlands: +00-31-33-45-15-000, New Zealand: Auckland: +64-9-976-6800; Wellington: +64-4-462-0780, Norway: +47 23 36 96 00, People's Republic Of China: Beijing: +86-10-6803-5588; Chengdu: +86-28-619-9333; Guangzhou: +86-20-8755-5900; Shanghai: +86-21-6466-1228; Hong Kong: +852-2202-6688, Poland: +48-22-8747800, Portugal: +351-21-4134000, Russia: +7-502-935-8411, Saudi Arabia: +9661 273 4567, Singapore: +65-6438-1888, Slovak Republic: +421-2-4342-94-85, South Africa: +27 11 256-6300, Spain: +34-91-596-9900, Sweden: +46-8-631-10-00, Switzerland: German: 41-1-908-90-00; French: 41-22-999-0444, Taiwan: +886-2-8732-9933, Thailand: +662-344-6888, Turkey: +90-212-335-22-00, United Arab Emirates: +9714-3366333, United Kingdom: +44-1-276-20444, United States: +1-800-555-9SUN OR +1-650-960-1300 Venezuela: +58-2-905-3800, Or Online at sun.com/store

Printed in USA 04/03 White Paper