

COTS Tenders and Integration Requirements

Soren Lauesen
IT University, Glentevej 67
DK-2400 Copenhagen NV
slauesen@itu.dk

Abstract

When buying COTS-based software, the customer has to choose between what is available. The supplier may add some minor parts, but not everything the customer wants. This means that the customer cannot write down his requirements and expect that they can all be met. A scoring system is necessary rather than traditional mandatory requirements. Requirements for integrating the new COTS system with other systems are particularly hard because suppliers may integrate in different ways and with different other systems. A related problem is that once the new COTS system is purchased, the COTS supplier may have a de-facto monopoly. Only he can expand the system or integrate it with other systems.

Experience shows that customers fail to deal with these issues adequately. As an example they may believe that asking for open interfaces is sufficient to guard them against monopoly. In this paper we analyze the problems and show ways to deal with them. We illustrate the problems and solutions with real-life examples from Electronic Patient Recording systems (EPR).

1. Background

Software is increasingly bought as COTS - Commercial Off The Shelf. COTS products range from simple components (e.g. sort procedures or user interface components), through complex operating systems and middleware (e.g. CORBA-based), to tool packages (e.g. MS Office) and complex application packages (e.g. ERP systems such as SAP or BAAN).

The COTS term should not be taken too literally. In practice, the COTS supplier may not only deliver the software off the shelf, but also extend it in customer-specific ways, for instance integrate it with other systems. He may also offer consultancy, data conversion, user training and support. Morisio & Torchiano [16] give a comprehensive taxonomy of COTS products.

The ability for two products to cooperate is called *interoperability*, and requirements dealing with it are called interoperability requirements. When two products cooperate closely, we say that they are *inte-*

grated. There are many degrees of integration as we will see below.

In this paper we will look at cases where the supplier delivers one or more COTS application packages, typically including middleware. He adds wrappers and glue to integrate his products with the customer's legacy systems. He may be willing to extend his COTS products in ways that meet specific customer needs, if they also seem suitable for future customers.

COTS products may be acquired in many ways. Here are some examples:

Components for in-house development

The customer's IT-organization looks around for suitable COTS components, buys them and uses them in their own systems. Semancik & Conger [18] give an example where the customer (NASA) built an autonomous file and distribution system. Their approach was to gather information about potential components from the Internet, local experts, and peers. They got trial versions from the vendors and tried out the products. Gradually they developed selection criteria and narrowed down their search to a few vendors.

Components for complex products

The customer is a product developer. He integrates the COTS components into his own product. Helokunnas and Nyby [9] give an example where the customer (Nokia) finds suitable suppliers, negotiates detailed interface requirements with them, and establishes long-term partnerships.

Business applications, company customer

The customer is a large company who wants a business application. Examples are ERP systems and banking systems. The customer will usually ask the supplier to do the necessary integration with his legacy systems. The acquisition approach may be similar to the approaches above: The customer gradually develops selection criteria, tries out the systems, and narrows down the search. He may negotiate conditions and long-term relationships with the supplier.

COTS tender for public organizations

The customer is a public organization who wants an application such as a health care system or a payroll

system. In the EU, public acquisitions have to follow strict rules to protect against corruption. All suppliers must have equal opportunities. For contracts above a certain amount (around 150,000\$), this is achieved through a tender process where the customer specifies the requirements and the selection criteria. The suppliers submit their proposals. The customer chooses between the proposals - based on the selection criteria. Prices, requirements and other conditions cannot be negotiated once the tender process has started.

This situation is much more difficult. The supplier systems are so complex that it is not feasible to try them out prior to the tender process.

In this paper we look at solutions for the tender situation. In particular, we look at requirements for integrating the new system with legacy systems and systems that the customer may acquire later. Although the solutions were developed for the public tender situation, they are also useful for other kinds of COTS acquisition.

2. Research in the area

Much research has been done about COTS in general, particularly about the vendor side - how COTS products are developed. Much less has been done about COTS software acquisition in practice. Maiden & Ncube [15] report on the lessons learned in selecting a software tool. A key lesson was that selection criteria were an exploratory exercise - you cannot define the requirements based solely on perceived customer needs. Brownsword et al. [6] outline a process for COTS software acquisition. Albert & Brownsword [1] explain that COTS solutions comprise much more than the COTS product itself. They all stress the iterative narrow-down approach (risk-based) and suggest that it can be handled in a semi-structured way (RUP-based). However, the tender situations we discuss don't allow iteration of this kind.

Until recently little research was done about the customer's attempts to integrate COTS products. Bansler & Havn [3] and Boehm & Abts [5] report on some of the difficulties and lessons from COTS integration in practice. Boehm & Abts give the important advice of early inspection of the COTS products and their integration abilities, rather than trusting the supplier's claims (a risk-driven approach). They also advise to choose open architectures to better deal with future products. Balk & Kedia [2] discuss a specific project with integrating nine COTS products. The integration was done by the customer's IT staff, and all the COTS products had already been used by the customer as stand-alone systems. Uncertainties about their functionality were thus reduced, and the requirements were not contractually strict. The integration required around 3,700 work hours. Hornstein & Willoughby [10] report about acquisition problems in NASA and how history had created artificial distinc-

tions of component types. These publications are useful background information, but don't help us express the requirements.

Gorton & Liu [7] discuss acquisition of COTS middleware and warn of the vague vendor specifications of the products, and the lack of compatibility in spite of "open interfaces". They explain in detail about their narrow-down approach when helping clients to select a middleware product. Their approach is inspiring, but since we look at application acquisition, the middle-ware requirements are not directly useful.

Interoperability in general is covered well, for instance with the many works on open source and various schemes and protocols for integration and reuse of software modules. Guo [8] and Wileden & Kaplan [19] give good overviews of the technological solutions. Bao & Horowitz [4] and Saur et al. [17] show two typical approaches to integrating largely incompatible products. Yakimovich et al. [20] made a classification of the integration technologies and their compatibility, and applied it to typical COTS products. Although these publications don't mention requirements, they have been a source of inspiration for the solutions below.

3. Research method

Principles

The first research question is *what are the real problems in COTS tenders*. Everybody in the area says that there are many problems, but they are vague about what the problems are. It is difficult to conduct research here because we are dealing with very big systems and very complex organizations - at customer side as well as supplier side.

Questionnaires might be sent to many companies, but they work badly for several reasons: (1) The actors in the market don't agree on terminology, so the questions are usually misunderstood. (2) It is hard to reach the right respondents and make them interested. (3) Respondents tend to answer what they hope will be going on in the next project, rather than what actually happened in the previous. (4) The respondents often haven't realized what the problems really are, so they cannot write about them. As a typical example, actors in the area tend to say that the main problem is bad project management. Asked what kind of mistakes project management made, they become silent. The fact is that "bad project management" can cover anything from immature time estimation at the supplier side, turmoil in the customer organization, to misunderstanding technical details about a third-party product. The cure depends very much on what the detailed problem is, but the actors don't dig down to this level of detail.

Interviewing the actors and studying the actual project documents give more reliable results, and this is what I have done over the years in many projects. The difficulty here is to get in contact with the right people. The matters are usually sensitive, and people are very

skeptical about researcher's attitudes. It is a matter of creating trust, respecting confidentiality, and respecting that practitioners are professional, knowledgeable people with good ideas that make them survive in the market place.

Once some important problems are known, the next research question is *how can the problems be removed?* In the COTS tender area you find few answers in literature, as explained above. However, literature may inspire you to come up with solutions. Another source of inspiration is what seemed to work well in other companies or in similar situations. Finally, analyzing the problems and trying to understand their causes, is an important source of inspiration. All of these sources have been important in this case, but the major source is innovation - ideas for new or modified ways of dealing with the problems. Afterwards, it is hard to explain how the ideas came up.

The final research question is *does the solution work?* This is the hardest one for several reasons. You cannot readily try out the solution in a real project. Under ideal circumstances it would take around a year, but usually the same customer will not do something similar for several years. In practice you have to be lucky to find another customer and another project where you can try out the ideas. Comparing the tender results with and without the new solution is very hard since so many other factors have changed at the same time.

Experiments with test groups (e.g. university students) is popular in research, but it works poorly for complex situations like these. When tried, practitioners tend to be very skeptical about the validity of the results.

The only way to get a fast indication whether a solution will work, seems to be expert judgement. Ask actors with deep experience in the area about the potential use of the solution. In the present case, this is how the proposed solutions were validated in the first place.

The approach outlined above is somewhat similar to *action research*: get involved in an actual project, find out what is really going on and what the actors actually do, suggest a solution and try it out in another project. The difference between action research and this project is that here we base the solutions on studies of several companies and several projects. Further, preliminary validation uses expert judgement, while the final validation takes years.

Practice

The basis for the present work was studies over a couple of years of five COTS tenders. The first was an EPR system for a shipyard, the next two were hospital systems for payroll and roster planning, the next a hospital system for patient administration, and the

last one was a meter reading and billing system for a municipality that supplied power, water, etc.

The main focus in these studies was how to deal with requirements to the user interface. At that time this was the most problematic area. As part of the first hospital studies we developed the Task & Support approach for requirements. The idea was to describe the user tasks without going into detail with who does what (similar to use cases, but without being explicit about the actors). The requirement is then to support these user tasks. The supplier describes for each task and subtask how he will support the user. The customer compares these descriptions, rates the solutions, and selects the winner.

The approach was first validated by expert judgement with one of the hospitals and with three suppliers. More than a year later, the hospital used the approach in a tender for patient administration. They reported excellent results and a reduction of tender costs to around one fifth. The approach and the results have been published elsewhere (Lauesen [11, 12]).

The last of the five studies (power and water supply) was intended to validate the Task and Support approach in another context. This project was not as successful for many reasons, one of them being an excessive user involvement (Lauesen & Vium [13]). However, it became apparent that the most problematic area had now shifted to integration requirements. In retrospect and when talking to suppliers, it was obvious that integration requirements had always been problematic, but now a much larger part of the project was about integration.

At first I had great difficulties coming up with reasonable integration requirements for a COTS tender. For some time I thought it was impossible. However, in cooperation with Vium, the customer, and later with one of the suppliers, I gradually developed the first outline of the solution presented in section 5. It was not really validated with experts at this point in time.

At the end of the study, I was fortunate to be involved as a consultant in two tender processes for EPR (Electronic Patient Recording). Here, the integration requirements were an even larger issue than with the power and water system. I had to elaborate the solution, and got some feedback from the hospital teams. They said that it looked okay. However, they were not really involved in a constructive way. Their focus was much more on the user interface aspect, user involvement (excessive), and the national EPR strategy. One of the hospitals had already purchased an integration platform and wanted to pursue the integration issue on their own. I continued the cooperation with the other hospital, however.

Since the experience from the power and water supplier was that integration *was* a critical issue, I visited the potential EPR suppliers (five of them) and talked to the expert groups that most likely would be involved in the tender. I showed them the outline of the integration requirements and asked whether it would be possible for

them to reply to these requirements, whether the requirements seemed fair, and whether they allowed them to describe the advantages of their own solution.

They had several comments that caused some additional changes to the integration requirements. This is discussed further in section 5.3. The hospital just took the integration requirements and patched them into a lot of other requirements that their IT department had made. As explained in section 5.2, this caused internal inconsistencies. And off the requirements went to the official bidding! It will take about half a year until it is possible to study the results. The inconsistent requirements are probably no real problem. Suppliers have much experience dealing with such things.

4. Problems in COTS tenders

In this section we will look closer at the integration requirements. We will illustrate the problems with real-life examples from acquisitions of EPR systems (Electronic Patient Recording systems). An EPR system records data about patients, their diseases (diagnoses), the treatments (patient services), and notes about why things are done and what the plans are. It must be possible to transfer data electronically to other health organizations or to statistical databases. Traditional patient records are largely text. Structured data such as measurements and diagnoses are often embedded in the text, and for this reason hard to use for overviews of the patient's state or for statistics. A good EPR system must help the clinical staff record all data in a more structured way and be able to show the data in many ways.

One of the big issues is that there are thousands of different patient services. They include laboratory tests, medicine prescriptions, preparation for surgery, X-ray pictures, CT scanning, food for the patient, psychological services, and so on. Each type of service may have its own data format. From a computer science point-of-view we are dealing with a class, *service*, that has thousands of sub-classes.

Furthermore, many services are handled by separate sub-systems (*production systems*) more or less automatically. As an example, some laboratory tests are carried out fully automated. Batches of sample glasses with bar codes are inserted in the analyzer, and a bit later the test results are available in the production system's database. A good EPR system must integrate its own database and functionality with a score of specialized production systems. These systems don't even have a common communication standard.

4.1. Difficult to explore existing products

Large COTS applications are too complex to try out in an exploratory manner. Typically the supplier would spend weeks setting up the system, and the cus-

tomers cannot do it himself. The customer might work closely together with the potential suppliers one by one to study their products, but this too is time consuming. There is also a risk that the customer bases his requirements on one particular system that he studied closely. This will unintentionally exclude other suppliers. Although the customer is aware of this problem, he cannot figure out how to phrase his requirements in a supplier-independent way, and in the public area it easily leads to accusations of not treating the suppliers equally.

The customer may study how successful the products are in other organizations, and many customers actually do this. However, the products develop so fast that this may give a false picture of what is available now. The result is that the exploratory narrow-down approach suggested in literature doesn't work well in the tender case. Even if it did, the customer would still have to go through a tender process and justify the final selection of the winner through compliance with requirements and selection criteria.

4.2. Mandatory requirements don't work

To run a tender process, the customer must elicit requirements in the traditional fashion. However, he cannot express the requirements in the traditional hard way, such as

R1: The EPR system shall not store the results in its own database, but retrieve them from Labsystem X when needed.

The problem is that this may be feasible for some EPR products, but for other EPR products it may either be impossible or cause very long response times. Such a requirement would unintentionally exclude a lot of suppliers.

Some analysts rightly say that R1 is a solution - not a requirement. To find the true requirement, we can ask the customer why he wants this solution. We can then replace R1 with this requirement:

R2: The EPR system must ensure that there is no difference between the shown results and the data in Labsystem X.

This requirement is unnecessary strict however, and it unnecessarily limits the possible solutions. In general it is impossible to find a supplier who can meet all the requirements that the customer dreams up. The customer wants to select the supplier that is closest to his wishes. The solution is to use *open-target requirements*, where the customer specifies his demands and expectations, while the supplier specifies how he can meet the expectations. Aware of the open-target approach, some customers specify their demands in this way:

R3: The system should share data with Labsys X. The customer expects that the latest results are always shown in the EPR system. The supplier is asked to explain his solution.

An open-target requirement doesn't arbitrarily exclude any suppliers but allow them to explain their solution. One supplier may reply that he always retrieves the data from Labsys X, another that he keeps a copy but updates it on a daily basis or at user request. Some suppliers may even say that they don't integrate with Labsys X at all, but they suggest another lab system that they integrate with.

Is it possible to verify open-target requirements? No, not in the traditional sense where a requirement is either met or not. Instead, the customer will give each supplier a score for how well he meets R3. He will give these scores when he assesses the proposals. The scores are the basis for selecting the winner.

In principle this works well, but in practice the supplier explanations may be long sales talks that are hard to compare. Below we will show a more structured approach that eases the comparison.

4.3. Product scope is fuzzy

The customer may not be sure how large a system he wants. Should it include middleware or should it use the platform that the customer has already? In the EPR case, should it include an X-ray system or should he get one from a third party? The best choice varies from one COTS system to another. Somehow the requirements must allow the suppliers to offer these things or integrate with a third-party product.

Some customers believe they can cut the entire system into pieces and purchase the pieces one by one. As an example, one hospital first purchased a middleware system. Next they wanted to purchase a series of smaller systems one by one that had to run on this middleware, for instance an X-ray system and a medicine system. Finally they planned to purchase the most important system - the EPR system - which would bind the other systems together from a user perspective.

Tempting as this may seem, experience shows that applications and middleware are so closely related that a system cannot easily be ported from one middleware system to another - even if they all claim to follow an open standard such as CORBA and J2EE (Gorton & Liu [7]; Liu & Gorton [14]). As a result, the customer may soon end up in a situation where no COTS product is available as the crucial EPR system. Selecting and buying middleware first, is a good strategy if the rest is developed from scratch. But if the rest is to be COTS products, the approach doesn't work well.

Another problem with this approach is that the customer arbitrarily has divided the entire system into

pieces ("modules"). He has defined the pieces from how his existing system happens to be organized. Modern vendors may not divide the world in the same fashion. They may find it awkward to deliver the requested pieces. Hornstein & Willoughby [10] report about similar problems in NASA, where history had created arbitrary splits of system functionality.

The general solution is to allow the supplier to deliver as much as possible - in particular the complex parts. This vastly reduces the integration problems. The right choice in the EPR case would be to acquire a middleware product combined with the crucial EPR system, but ensure that third parties can add new production systems later. Below we will show requirements that specify this.

4.4. High-risk areas handled too late

Once the contract is signed, the supplier should just deliver as promised. There may be a long period of time where the supplier develops glue-ware and extends his system, and at the end the customer will use the system. This is the ideal. However, it often happens that the supplier cannot fulfill his promises no matter how hard he tries. Typical trouble areas are performance (response time) and integration with other products. As an example, the supplier's present installations have around 30 users and perform well with them. But when facing the planned 2000 users, the response time skyrockets. Or the integration to system X sounded easy, but in practice it didn't work as expected.

Up front most customers don't care about this. What is the problem, they ask. If the supplier doesn't deliver as promised, he won't get his money; he may even have to pay a penalty. These customers assume that the supplier is just lazy and money will cure him. The fact may be that the supplier is unable to solve the problems. Such projects can drag on for years. The customer never gets an adequate system and the supplier loses fortunes. Court cases don't cure anything.

The root problem here is that the parties too late face the high-risk areas. The supplier tends to delay the hard parts and deliver the easy ones first. Ideally, the customer should ask for proof of the high-risk areas before selecting the winner. Often this is expensive, and it is unrealistic that all the proposers can do it at proposal time.

4.5. The COTS supplier gets a monopoly

Once the customer has got a complex system, he may want to extend it in various ways or integrate it with new systems. At this point the COTS supplier may have a monopoly. Only he can extend the COTS system. Some suppliers actually exploit this situation and charge unfair prices for extensions.

If COTS was a truly open market, the customer could just find another vendor - in the same way as you can switch from Ford to Toyota. Unfortunately, it is not so

15. Laboratory system

The customer expects integration with his present system, Labsystem X, but may consider changing to a new lab system. The technical interfaces to Labsystem X are specified in Appendix ...

Degree of integration:	Suggested or offered solution:
1. The vendor offers an alternative system that meets the functional requirements in section ...	
2. The vendor offers integration with Labsystem X as specified in points 3 to 10 below.	
3. The user starts Labsystem X through the EPR system, then logs into Labsystem X, specifies the patient ID and requests the lab service through X's screens.	
4. As possibility 3, but the user doesn't have to log in and specify the patient ID.	The user and patient ID are transferred automatically.
5. The user requests lab services through the EPR system's screens in the same way as for other services.	The EPR system uses the API interfaces to Labsystem X.
6. States and results of the lab service are visible in the EPR screens in the same way as for other services.	
7. The EPR system can warn users about pending lab results in the same way as for other services.	
8. The EPR system shares database with Labsystem X. In this way EPR data and lab data are always identical.	
9. Lab data is periodically transferred from Labsystem X to the EPR system (replicated databases).	
10. Data for a single patient is transferred at user request from Labsystem X.	

Figure 1. Section 15 of an EPR requirements specification

Specifies optional integration with an existing customer system and an optional inclusion of a different system.

easy with complex software applications. It often takes a couple of years to replace an application - primarily due to organizational changes.

Customers try to prevent the monopoly by asking for "open interfaces" to the COTS product. Typically, the supplier will reply: "Yes, our system has open interfaces - it uses CORBA". Much later the customer realizes that although CORBA is used, the detailed formats and meanings of messages and API-calls are not described, and the supplier considers this information proprietary. In order to prevent the monopoly, more than "open interfaces" are needed. Below we will show requirements that better guard against monopoly.

5. Dealing with the problems

In this section we will show solutions to the problems above. We will again illustrate the principles with an EPR acquisition.

5.1. Optional sub-products

When the product scope is open, we need a way to allow the supplier to influence the scope. The solution is to use open-target requirements and let the supplier specify which of them he will meet and how. It is

important to state the requirements as user demands, rather than in technical terms. The user view better allows the customer to assess the consequences of the supplier's solution.

Figure 1 shows section 15 of an EPR requirements specification. It is a set of user-oriented requirements that specify to what extent the customer's existing laboratory system will be integrated with the new EPR system. As an example, requirement 1 allows the supplier to offer an alternative lab system. The supplier writes his reply in column two. He may say "No, I don't offer an alternative system" or he may say "yes" and give a price for this option. Requirement 2 allows him to offer integration with the existing lab system. There are thus four basic options; the supplier offers only an alternative system; he offers only an integration; he offers both (and the customer can choose); or he doesn't deal with the lab system (a third party who knows the existing lab system might do the integration).

If the supplier delivers an alternative lab system, he delivers a larger product than if he just integrates with the old one. However, the integration may be so costly to him, that the larger product is cheaper than the integration.

5.2. Degrees of product integration

Figure 1 is also an example of how to avoid the mandatory requirements by means of open-target requirements. The supplier is not asked to do the integration in a specific way. He is asked to specify in a structured way how close the integration will be. Requirement 3 is the weakest kind of integration. The supplier just allows the user to switch over to the lab system, but doesn't interfere in any other way. Requirements 6 and 7 state that from the user's point of view, the lab services look like any other patient service. Requirements 8 to 10 allow the supplier to specify the actuality of data. When he fills in column two, he may go into more detail with the solution.

The figure shows two examples where the customer has pre-filled the fields in the second column. He has stated a solution he suggests, but not as a requirement. The supplier may replace it with the solution he offers. Experience is that these suggestions help the supplier understand what the customer expects. They may also allow the supplier to explain that his product exceeds the customer's expectation.

Which degrees of integration should we describe in the table? The example deals with what is important in this project. Guo [8] and Yakimovich et al. [20] have more generic lists of integration dimensions, and our choice here was inspired by their list. Gorton & Liu [7] explain about criteria for middleware products. They refer to their 150 item list of requirements to consider. It is proprietary, however. Such lists are precious extracts of long experience!

What if the supplier has a solution that we didn't know or didn't ask about? This is not a serious problem. The supplier just explains his solution in a right-hand box. Comparison is a bit harder - that is all.

Creating conflicting requirements

These product integration requirements were used by the hospital in a tender. They just trusted that they were okay and patched them into a lot of requirements made independently by the IT department. However, they didn't notice that the old R2 requirement was still there (with a new number, of course):

(R2): The EPR system must ensure that there is no difference between the shown results and the data in Labsystem X.

This requirement is in conflict with the softer integration requirements. It doesn't for instance allow periodic transfer of data as in requirement 9 above. The suppliers will hopefully point this out to the customer.

5.3. Integrating with future products

In the future, the customer may want to integrate the COTS-based system with other systems. In order

to avoid that the supplier gets a monopoly for the integration, the customer should require that a third party can do the integration. Two kinds of requirements are involved: the interfaces to be provided, and the human ability to do the integration. Figure 2 shows how this can be handled in an EPR tender.

The example talks only about integrating with new "production systems", but the requirements show that the concept of a "production system" is very wide and would allow many kinds of systems. Requirement 1, for instance, says that the new system can retrieve and update all data specified for the EPR system. Many systems can be integrated in this way. The supplier will in column two explain to what extent he can meet this requirement.

Requirement 2 specifies EPR functionality that must be available to the new production system. Very little is required because most of the communication is done by means of the data exchanged according to requirement 1. (Some more is required in the real case, for instance exchange of security information, but we have omitted it here.)

Requirements 1 and 2 specify what the EPR system should do in its role as a server for the new production system. Requirements 3 and 4 specify what the EPR system should do in its role as a client that requests services from the new production system.

One thing is that the two systems in principle must be able to exchange data and invoke functionality from each other. Another thing is whether a third party is able to make the systems do it. Requirements 5 through 8 aim at this.

Requirement 5 says that the technical interfaces must be documented. A requirement for documentation is very common, but it is also very common that the documentation actually supplied is useless. How can the customer guard against this? Requirement 5 says that it must be understandable and adequate for a software house. This requirement will be verified by having a software house check this. The software house will also check that the documentation matches what the system actually does. In order to reduce the risk, the supplier is asked to supply sample documentation as part of the proposal.

Some customers have experienced that although the technical interfaces are documented, neither customer, nor third parties are allowed to use the interfaces. In cases where the supplier also operates the system, he may even claim a right to the data stored on behalf of the customer. Requirements 6 and 7 guard against this. (In many cases these requirements are considered part of the contract rather than part of the requirements. This is not important as long as they are somewhere.)

Finally, requirement 8 asks for the response time to be specified for each of the functions on the technical interfaces.

16. New production systems

The customer expects that new production systems can be integrated with the EPR system by a third party. A new production system might for instance be another laboratory system, a booking system, or an expert system that uses clinical data about a patient for searching international medical databases.

A production system may be tailor-made or a COTS system where third party adds glue code or adapters for integration with the EPR system.

Requirements:	Suggested or offered solution:
Interfaces - EPR in the server role:	
1. The new production system can retrieve and update data in the EPR system. This data should include data described in section ...	Might be done with messages, an API call or an SQL-query.
2. The new production system can invoke functionality in the EPR system. The functionality includes reporting warning events and printing results on department printers monitored by the EPR system.	
Interfaces - EPR in the client role:	
3. The new production system can allow the EPR system to retrieve and update production data.	Might be offered as messages, an API call or an SQL-query.
4. The new production system can provide functionality for the EPR system. The functionality includes requests for service, warnings about pending requests, and print of requests.	
Documentation and rights:	
5. The technical interfaces to the EPR system (e.g. messages and API calls) and the format of data that can be retrieved or updated in the EPR system must be documented. The documentation must be understandable to a third-party software house and found suited for the intended development purpose.	The vendor should specify how much training is needed to understand the documentation. Sample documentation should be submitted as part of the proposal.
6. The customer and third parties must have the right to use the documentation and the technical interfaces.	
7. Third parties must have the right to extract and use data from the EPR system with the permission of the customer.	
8. Response times for the various interface functions should be specified, including their dependencies of database sizes and hardware platforms.	

Figure 2. Requirements specifying to what extent a third party can extend the system.

Expert judgement

The first outlines of the integration requirements were presented and discussed with five potential suppliers, one by one. We will explain some of the more interesting comments.

Suppliers usually strongly resist that third parties integrate with their product. They find that the technical stability of their system is at stake (and maybe also their monopoly). What was the reaction during the presentation? It turned out that all suppliers accepted the need for third parties. Some of them even said that they knew they had refused earlier, but things had changed. They had to be competitive and open up.

Generally, they found the requirements above suitable for the purpose. One supplier had a system with an open interface that however was very hard to learn. He insisted that some weeks be allowed for the third party to learn the interface. Another supplier ex-

plained that any programmer would be able to use their open interface right away. The result was that we in the right-hand box for requirement 5 asked the supplier to specify how much training was needed.

Requirement 8 was more of a problem. The suppliers understood the need for something like this. If the figures are unknown, it is impossible to assess what the interfaces can be used for in practice. However, they would hate to reply to it (some of them probably wasn't sure they would be able to do so). We kept the requirement, but as for all other requirements, the target is open and the supplier may reply that he doesn't have the figures. A supplier who can supply the figures will score higher than one who doesn't, but it is not sure he will win for that reason.

Selection criteria

1. Efficient support for the basic clinical tasks.
2. Adequate response times for the specified number of transactions and users.
3. Ability for the customer and third parties to extend the EPR system.
4. Integration with the customer's existing systems.

The customer considers these criteria the most important for long-time success of the system. They are also considered the most risky areas. Failure in one of these areas will be very hard to repair late in the project.

The supplier should as far as possible demonstrate how far the criteria are met in the present system. This may be done during the presentation meeting when the customer evaluates the proposal. If it cannot be done at this time, the criteria must be demonstrated during the first six weeks after signing the contract.

Figure 3. The selection criteria from the tender documents.

5.4. Step-wise narrow-down and a trial period

While the criteria for *component* acquisition can be modified as new insights are obtained, this is not possible in public tenders. A tender approach narrows down the suppliers according to predefined steps and criteria. A typical tender process proceeds in this way:

Pre-qualification. The customer announces that a tender will be made in a certain area, for instance health care. No requirements are announced at this point, but the selection criteria for the pre-qualification are stated. They will typically be: *The supplier's financial status and track record, reference customers, experience in the application area (for instance health care)*. Suppliers are invited to apply for pre-qualification.

The customer selects a limited number of suppliers, typically five. The effort for suppliers as well as customer is modest compared to the effort of replying to the real requirements.

Tender. Around a month later, the customer sends the requirements specification, the final selection criteria and other tender material to the pre-selected suppliers. The suppliers get between 26 and 40 days to send a proposal. The customer selects the winner from an assessment of the proposals. As part of this, each supplier has a presentation session where he can explain and demonstrate his proposal.

At this point the traditional narrow-down process is finished. In many ways it works all right. However, the supplier has only *explained* what he will deliver,

but he has not *proved* that he can do it. To reduce the risks for both parties, the supplier should deal with the high-risk areas very early.

Trial period. The solution is to encourage the winner to prove already at proposal time that he can do as promised. For areas where it would be too costly for him to prove it at this time, he is given a short period to do it after signing the contract. As an example, he may prove that he can meet the response time requirements by setting up an environment where 2000 users are simulated. Although this is not the final verification of the response time requirements, it is sufficient evidence at this point. Notice that such a trial period makes sense for COTS-based systems, which largely exist already. It would not be realistic for systems developed from scratch.

In one case the customer selected two winners and paid both of them around 100,000\$ to go through the trial period. At the end of the period, he selected one of them to do the rest of the project. This is another step in the narrow-down process. For each step more work is needed to reach the next step, but fewer suppliers have to do it.

Selection criteria. Which criteria should be used for the selection? In principle the answer is easy: look at the major risks. A major risk is something that can wreck the project completely. Examples are inability to provide adequate response times, or inability to extend the system without the supplier having a monopoly. Select suppliers that provably eliminate these major risks. Minor risks, such as some missing functionality, can be overcome by the supplier - at least if money is at stake - and need not enter the selection criteria.

Figure 3 shows an example of how the selection criteria and the trial period can be expressed in an EPR tender. Each criterion is a summary of many requirements. As an example, criterion 1 (*Efficient support for the basic clinical tasks*) comprises the scores for how well the EPR system supports a few basic user tasks. Criterion 3 (*Ability for the customer and third party to extend the EPR system*) refers to the future-product requirements in Figure 2, plus additional requirements for the ability to add new types of services and associated screen pictures.

Discarding the best supplier too early. A step-wise selection of this kind has the disadvantage that the best supplier - everything considered - may have been discarded early in the process. As an example, a new supplier entering the application domain with a superior product, may be discarded already in the pre-qualification stage. While the more exploratory narrow-down approaches could deal with this situation, the tender process cannot.

6. Conclusion

Buying COTS products that integrate with other products is a difficult business. Based on practical experience from tender processes, the paper suggests these improvements of current practice:

1. Use open-target requirements that allow the supplier to explain his solution in a semi-structured way.
2. Express each requirement as a user demand rather than a technical feature in order to better see the consequences of the supplier's solution.
3. Specify degrees of integration and ask the supplier which degrees he can provide.
4. Specify requirements that ensure that a third party can extend the product. The requirements must cover the interfaces needed, the usefulness of the documentation, and the right to use the documentation and interfaces.
5. Use a short trial period immediately after signing the contract to get a proof that the supplier can handle the high-risk areas of the project.

Acknowledgements

I want to thank Jens-Peder Vium, Henrik Willumsen, David Simonsen, Niels R. Larsen and Henrik Lindholm for some of the ideas behind this paper, Andrew Gabb for assuring me that I was dealing with important problems, and Torben Elin for comparing the approach with what he - as a customer - had been trying to do.

7. References

- [1] Albert, C. & Brownsword, L.: Meeting the challenges of Commercial-Off-The-Shelf (COTS) products. In: J. Dean & A. Gravel (Eds.): International Conference on COTS-Based Software Systems, ICCBSS 2002, LNCS 2255, pp. 10-20.
- [2] Balk, L.D. & Kedia, A.: PPT: A COTS integration case study. International Conference on Software Engineering, ICSE 2000, pp. 42-49.
- [3] Bansler, J.P. & Havn, E.C. (1994): Information systems development with generic systems. In: Walter R.J. Baets (ed.), Second European Conference on Information Systems. Nijenrode University Press, 1999, pp. 707-715.
- [4] Bao, Y. & Horowitz, E.: Integrating through user interface: A flexible integration framework for third-party software. Proceedings of COMPSAC '96. IEEE Computer, 1996, pp. 336-342.
- [5] Boehm, B. & Abts, C.: COTS integration: Plug and Pray? IEEE Computer, January 1999, pp. 135-38.
- [6] Brownsword, L., Oberndorf, T. & Sledge, C.A.: Developing new processes for COTS-based systems. IEEE Software, July/August, 2000, pp. 48-55.
- [7] Gorton, I. & Liu, A.: Streamlining the acquisition process for large-scale COTS middleware components. In: J. Dean & A. Gravel (Eds.): International Conference on COTS-Based Software Systems, ICCBSS 2002, LNCS 2255, 2000, pp. 122-131.
- [8] Guo, Jiang: Interoperability technology assessment. Elsevier Science, Electronic notes vol. 65 No. 4., 2000.
- [9] Helokunnas, T. & Nyby, M.: Collaboration between a COTS integrator and a vendor. In: J. Kontio & R. Conradi (Eds.): European Conference on Software Quality, ECSQ 2002, LNCS 2349, pp. 267-273.
- [10] Hornstein, R.S. & Willoughby, J.K.: Realizing the potential for COTS utilization: a work in progress. In: J. Dean & A. Gravel (Eds.): International Conference on COTS-Based Software Systems, ICCBSS 2002, LNCS 2255, 2002, pp. 142-150.
- [11] Lauesen, S.: Software requirements - styles and techniques. Addison-Wesley, 2002.
- [12] Lauesen, S.: Task Descriptions as Functional Requirements. IEEE Software 2003, March/April, pp. 58-65.
- [13] Lauesen, S. and Vium, J.P.: Experiences from a tender process. International Workshop on Requirements Engineering, REFSQ'04, 2004.
- [14] Liu, A. & Gorton, I.: Accelerating COTS middleware acquisition: The i-Mate process. IEEE Software, March/April 2003, pp. 72-79.
- [15] Maiden, N.A. & Ncube, C.: Acquiring COTS software selection requirements. IEEE Software, March/April, 1998, pp. 46-56.
- [16] Morisio, M. & Torchiano, M.: Definition and classification of COTS: a proposal. In: J. Dean & A. Gravel (Eds.): International Conference on COTS-Based Software Systems, ICCBSS 2002, LNCS 2255, pp. 165-175.
- [17] Saur, L.D., Clay, R.L. & Armstrong, R.: Meta-component architecture for software interoperability. IEEE Computer, November, 2000, pp. 75-84.
- [18] Semancik, S.K. & Conger, A.M.: The standard autonomous file server, a customized, off-the-shelf success story. In: J. Dean & A. Gravel (Eds.): International Conference on COTS-Based Software Systems, ICCBSS 2002, LNCS 2255, pp. 234-244.
- [19] Wileden, J.C. & Kaplan, A.: Software interoperability: Principles and practice. In: Proceedings of Software Engineering 1999, ACM, pp. 675-676.
- [20] Yakimovich, D., Bieman, J.M. & Basili, V.R.: Software architecture classification for estimating the cost of COTS integration. In: International Conference on Software Engineering, ICSE '99, ACM, pp. 296-302.