

*System Name*  
Software Architecture Description

*Author Name*  
*Contact Details*  
*Version*  
*Date*

## 1. Version History

<b><i>Version</i></b>	<b><i>Date</i></b>	<b><i>Author</i></b>	<b><i>Comments</i></b>
1	July 08	Eoin Woods and Nick Rozanski	Initial version of template
2	November 11	Eoin Woods and Nick Rozanski	Revised for Edition 2

## 2. Table of Contents

1. Version History .....	2
2. Table of Contents .....	3
3. <a href="#"><u>Instructions</u></a> .....	4
4. Introduction .....	5
4.1 Purpose and Scope .....	5
4.2 Audience .....	5
4.3 Status .....	5
4.4 Architectural Design Approach .....	5
5. Glossary .....	6
6. System Stakeholders and Requirements .....	7
6.1 Stakeholders .....	7
6.2 Overview of Requirements .....	7
6.3 System Scenarios .....	7
7. Architectural Forces .....	9
7.1 Goals .....	9
7.2 Constraints .....	9
7.3 Architectural Principles .....	9
8. Architectural Views .....	10
8.1 Context View .....	10
8.2 Functional View .....	11
8.3 Information View .....	13
8.4 Concurrency View .....	16
8.5 Deployment View .....	18
8.6 Development View .....	20
8.7 Operational View .....	22
9. System Qualities .....	23
9.1 Performance and Scalability .....	23
9.2 Security .....	23
9.3 Availability and Resilience .....	23
9.4 Evolution .....	23
9.5 Other Qualities .....	24
10. Appendices .....	25
10.1 Appendix: Decisions and Alternatives .....	25
10.2 Appendix: Questions and Answers .....	25
10.3 Appendix: References .....	25

### **3. Instructions**

This template provides an outline for a view-based software architecture description document, of the style suggested by the approach in the book *Software Systems Architecture*, 2<sup>nd</sup> Edition by Nick Rozanski and Eoin Woods (Addison Wesley, 2011).

The text in this style is all instruction text, which you should remove as you write the document. Once you have done this, the template is essentially a set of headings and some pro-forma text.

This template is a superset of the content of a typical AD. If there are sections of the template that are not relevant to your project then you should probably omit them. Similarly if there are topics that are important to your system but not part of the template, then you should add appropriate sections to it.

For more information about the architectural approach underpinning the template, the concepts presented and the techniques used to create effective architectural descriptions, see the web site <http://www.viewpoints-and-perspectives.info>.

## **4. Introduction**

### **4.1 Purpose and Scope**

Explain the purpose and scope of the document.

Primarily this is to document the architecture for the stakeholders, to ensure that it meets their goals and concerns and that the proposed architecture is correct, complete and fit for purpose.

While you should avoid presenting a lot of material available elsewhere, it may also be useful to do some or all of the following in the AD:

- summarise the project context, goals and objectives
- confirm scope and exclusions
- present an overview of goals and drivers, requirements etc
- record important decisions made and their rationale
- present alternatives considered and their reasons for rejection
- bring together other important information not captured elsewhere

### **4.2 Audience**

Define the audience of the document.

Probably the most important audience are the developers who will be designing and building the system, along with the sponsor who will be paying for it. However there are a number of other stakeholders who have an interest in the AD, as listed in the System Stakeholders and Requirements chapter. And of course the architect is a stakeholder as well.

### **4.3 Status**

Explain the current status of the architecture and of this architectural description.

Is it still in progress? Being implemented? In production? You may also want to describe future plans for the document (eg will be reissued as Definitive after comments received by stakeholders).

### **4.4 Architectural Design Approach**

Explain the overall architectural approach used to describe and develop the content of the document (e.g. explain viewpoints, views and perspectives). If necessary explain the architectural views that you're using and why each is used.

## 5. Glossary

Define any terms, acronyms or abbreviations that might be unfamiliar to the target audience. This should include both business terms and technology / architectural terms.

If the glossary is long, create a separate document and reference it here.

Term	Definition
term	definition of the term

## 6. System Stakeholders and Requirements

### 6.1 Stakeholders

Define each of the key stakeholders and stakeholder groups, explaining their interest, needs and concerns for the system.

A stakeholder is anyone who has an interest in or concern about in the system documented in the AD. Consider the following stakeholder groups.

- **Acquirers**, who pay for the system.
- **Assessors**, who check for compliance.
- **Communicators**, who create documents and training.
- **Developers**, who create the system.
- **Maintainers**, who evolve and fix the system.
- **Production Engineers**, who are responsible for the deployment environment.
- **Suppliers**, who provide parts of the system.
- **Support Staff**, who help people to use the system.
- **System Administrators**, who keep it running.
- **Testers**, who verify that it works.
- **Users**, who have to use the system directly.

And of course the architect is also a stakeholder in the AD.

### 6.2 Overview of Requirements

Summarise the key functional and quality property (non-functional) requirements for the system.

Functional requirements define what the system is required to do (for example, update customer name and address details). Quality properties (aka non-functional requirements) define how the system must behave at run-time or design time (for example, it must respond to requests within three seconds under a given load; it must be available 99.99% of the time; it must be possible to extend the system to meet certain types of new requirement without having to undertake a major redesign).

Avoid going into too much detail which is presented elsewhere; refer to external sources, such as requirements documents, SLAs, existing systems and so on, wherever possible. Requirements should be numbered so that you can refer to them unambiguously elsewhere.

Reference	Requirement Description
R1.	
R2.	

### 6.3 System Scenarios

List, and briefly outline the most important scenarios that matter to the key stakeholders and/or can be used to illustrate the system's ability to meet its most important requirements.

A scenario describes a situation that the system is likely to face in its production environment, along with the responses required of the system. You should consider both functional scenarios (things that the system must do usually in response to an external event or input) and system quality scenarios (how the system should react to a change in its environment, such as an increase in workload).

In most cases the scenarios take a significant amount of space and it is often appropriate to record them in a separate document to avoid the AD getting too large.

### 6.3.1 Functional Scenarios

Functional scenarios model things that the system must do response to an external stimulus (eg an event or input).

<b>Scenario Reference</b>	
<b>Overview</b>	
<b>System State</b>	
<b>System Environment</b>	
<b>External Stimulus</b>	
<b>Required System Response</b>	

### 6.3.2 System Quality Scenarios

System quality scenarios model how the system should react to a change in its environment (such as an increase in workload or a security breach).

<b>Scenario Reference</b>	
<b>Overview</b>	
<b>System Environment</b>	
<b>Environment Changes</b>	
<b>Required System Behaviour</b>	

## 7. Architectural Forces

### 7.1 Goals

List the main architectural goals and business drivers for the project.

A goal is something that the project wants to make happen (eg to simplify the customer management processes) while a business driver is some external force which shapes the project (eg the level of customer complaints is increasing).

Ask yourself questions like the following: What are the key targets that you are setting yourself as the architect? Do you aim to reuse existing software? Or develop the system at minimum cost? Or are you aiming for very high reliability? Or ... ?

### 7.2 Constraints

List the main architectural constraints that the project must respect.

A constraint is something that limits your architectural choices: for example, the project must be completed by Christmas, it must be implemented in Java, or it must conform to a certain operational model.

You should also refer to any specific standards or regulation which govern the architecture.

### 7.3 Architectural Principles

Explain the architectural design principles which have shaped the architecture.

A principle is a fundamental statement of belief, approach, or intent that guides the definition of the architecture. It may refer to current circumstances or a desired future state. A good principle is constructive, reasoned, well-articulated, testable and significant.

Each principle should be justified by a rationale and may be supplemented by some implications. For example, a principle on the use of open standards may have as its rationale the drive for interoperability, and as an implication the need to assess and agree the appropriate standards which apply to each component.

You may also choose to include more specific principles in the Views chapter.

<b>Principle Reference</b>	(unique number)
<b>Principle Statement</b>	(brief statement of the principle)
<b>Rationale</b>	
<b>Implications</b>	
<b>Further Information</b>	

## 8. Architectural Views

### 8.1 Context View

The context view of the system describes the relationships, dependencies and interactions between the system and its environment (the people, systems and external entities that it interacts with).

#### 8.1.1 Context Diagram

Use a context diagram (and supporting explanation) to explain the environment in which the system operates and the external entities it interacts with. Briefly define here each of the external entities and the important interactions that the system has with them.

The context diagram is usually presented as a simple, high-level picture which shows the system's boundaries and its adjacent external entities. The external entities are typically other systems, but may also be physical devices, "black box" external organisations, or more granular software components. Interactions may be data flows (interfaces) or control flows (eg invoking a service or exposed function).

You normally present the system itself in the diagram as a single box or component. External entities are also normally presented as single boxes or components, since you often don't know (or care) about their internal implementations.

An example context diagram is given below. This uses a "neutral" boxes-and-lines notation; you may want to use a more formal modelling language such as UML instead.

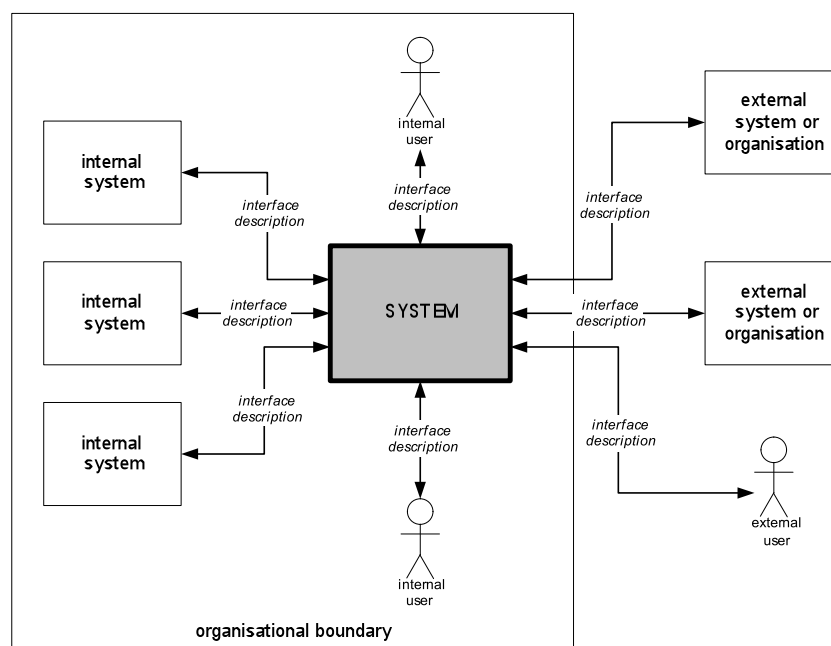


Figure 1. System Context

#### 8.1.2 Interaction Scenarios

Where you have complicated interactions between your system and external entities, consider modelling some of the expected interaction sequences using interaction scenarios. These can help to uncover implicit requirements and constraints (such as ordering, volume or timing constraints) and helps to provide a further, more detailed level of validation.

You can capture interaction sequences using UML sequence diagrams or bulleted lists of interactions.

## 8.2 Functional View

The functional view of the system defines the system's architecturally significant functional elements, the responsibilities of each, the interfaces they offer and the dependencies between elements.

Place a functional model here (e.g. a UML component diagram) and explain its content in the subsections below. A functional element is a well-defined part of the runtime system that has particular functional responsibilities and exposes interfaces that connect it to other functional elements.

Focus on the important functional elements in your architecture. In general you should not model the underlying infrastructure here unless it performs a functionally significant purpose (for example a message bus that links system elements and transforms data exchanged between them).

If your architecture is functionally complex you may choose to model it at a high level and then decompose some elements in further sub-models (functional decomposition).

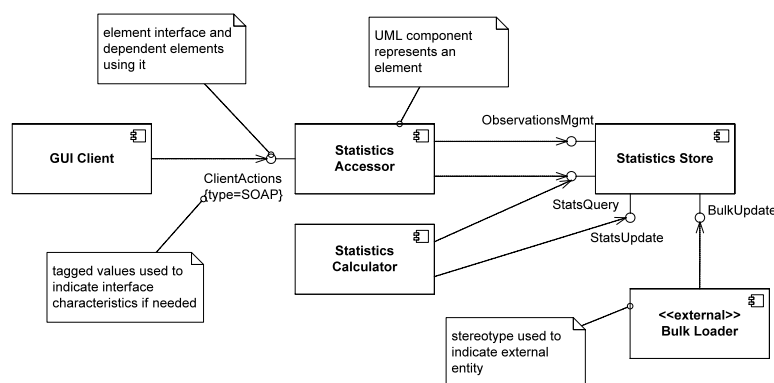


Figure 2. Functional Model

### 8.2.1 Functional Elements

Define the responsibilities and interfaces offered and/or required by each functional element. Alternatively, if you are using a modelling approach like UML you might choose to keep the main descriptions in the UML model repository and summarise the information here, referencing the model(s).

If you have used functional decomposition in the previous section, you can structure this section to align with your functional hierarchy.

Element Name	
Responsibilities	
Interfaces – Inbound	
Interfaces – Outbound	

### **8.2.2 Functional Scenarios**

Use one or more interaction diagrams to explain how the functional elements interact via their interfaces, in order to meet some of the key system functional scenarios.

### **8.2.3 System-Wide Processing**

Define how any system-wide processing will be handled (for example, if you have a message-oriented system, how will you deal with message delivery errors across the system).

## 8.3 Information View

The Information view of the system defines the structure of the system's stored and transient information (e.g. databases and message schemas) and how related aspects such as information ownership, flow, currency, latency and retention will be addressed.

### 8.3.1 Data Structure

Define or reference any architecturally significant data structures for stored and transient data, such as overview data models or message schemas.

At this level you should keep the number of entities small – no more than 20 or so if possible. It is not necessary to be 100% normalised – for the sake of clarity it is acceptable to have some many-to-many relationships for example. Don't try and illustrate every entity and relationship here or your readers will get lost in the detail.

It may also be useful to logically group entities together that are semantically related in some way – for example, all data related to customer name and address. This may help your readers to understand the data items and the relationships between them.

Here is an example data structure model which uses classic ERD notation. You can also use class diagrams here although that may be too granular a level of detail for an AD. An alternative, should you wish to use UML, is to illustrate the information structure at the package, rather than the class, level.

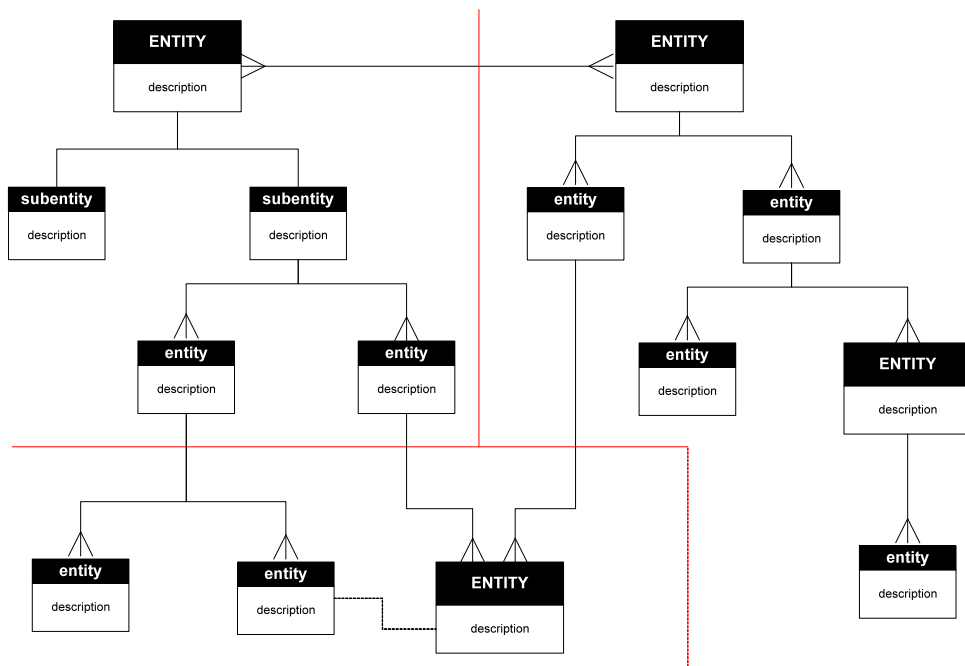


Figure 3. System Data Structure

### 8.3.2 Data Flow

If it is not clear from the functional view's interaction diagrams, define how data flows through the system from one component to another and to external components.

As with the data structure diagram, keep this simple and focus on no more than about 10-15 key functional elements. Don't try and illustrate every data flow here or your readers will get lost in the detail.

An example is shown below using a data flow diagram.

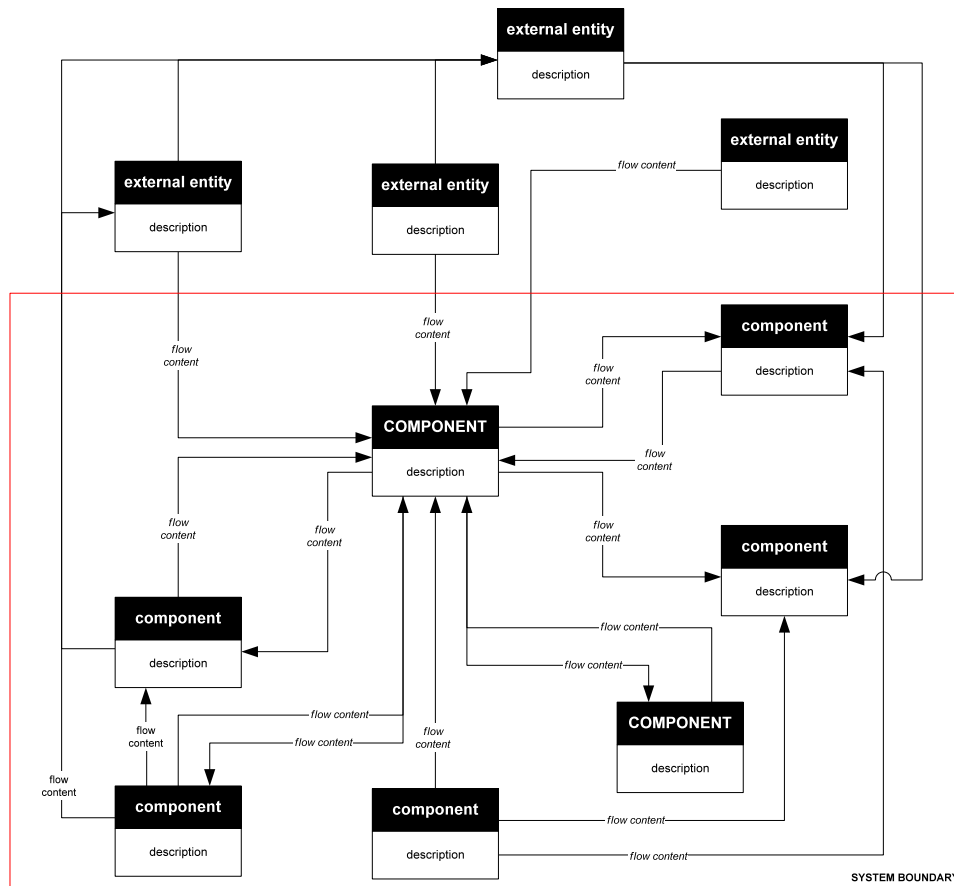


Figure 4. System Data Flow

### 8.3.3 Data Ownership

If data is owned by more than one entity or part of the system, define who owns which pieces of the data and explain how any resulting problems will be handled.

In the example below, it can be seen that there are issues with entity 4 which can be updated by System D which is not the owner. The AD should explain how this inconsistency will be managed.

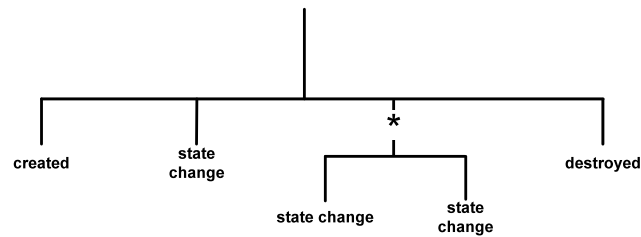
Entity	System A	System B	System C	System D
entity 1	MASTER	r/o copy	reader	reader
entity 2	reader	MASTER	none	reader
entity 3	none	reader	MASTER	reader
entity 4	MASTER	none	none	reader updater deleter

### 8.3.4 Information Lifecycles

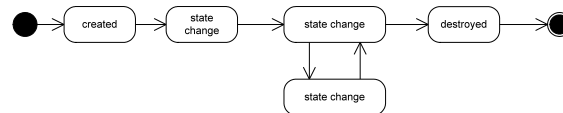
If key entities have complicated lifecycles then model the way that their state changes over time.

Focus on a few key entities whose transitions help to illuminate key features of the architecture, rather than just created/updated/updated/updated/destroyed.

There are two common techniques for modelling information lifecycles, entity life histories and state transition diagrams. Both are useful; choose one style and stick to it throughout the AD.



**Figure 5. Entity Life History**



**Figure 6. State Transition**

### 8.3.5 Timeliness and Latency

If information needs to be copied around the system or is updated regularly, explain how timeliness and latency requirements will be addressed.

### 8.3.6 Archive and Retention

Explain how will archive and retention requirements will be met by the system.

## 8.4 Concurrency View

The Concurrency view of the system defines the set of runtime system elements (such as operating system processes) into which the system's functional elements are packaged.

If the concurrency structure is complicated or it isn't obvious from the information in the other views, define how functional elements will be packaged into processes and threads and explain how they interact safely and reliably using suitable inter-process communication mechanisms. This can be achieved via a UML model (using stereotypes), by using a special purpose concurrency modelling language, or by creating an informal notation for the situation at hand.

### 8.4.1 Concurrency Model

Model the processes, process groups and threads, and the interprocess communication channels between them.

You may also choose to model the mechanisms used to protect the integrity of data and other resources shared between concurrent execution units, such as mutexes or semaphores.

You can use a UML component model to represent the information graphically, stereotyping the components appropriately.

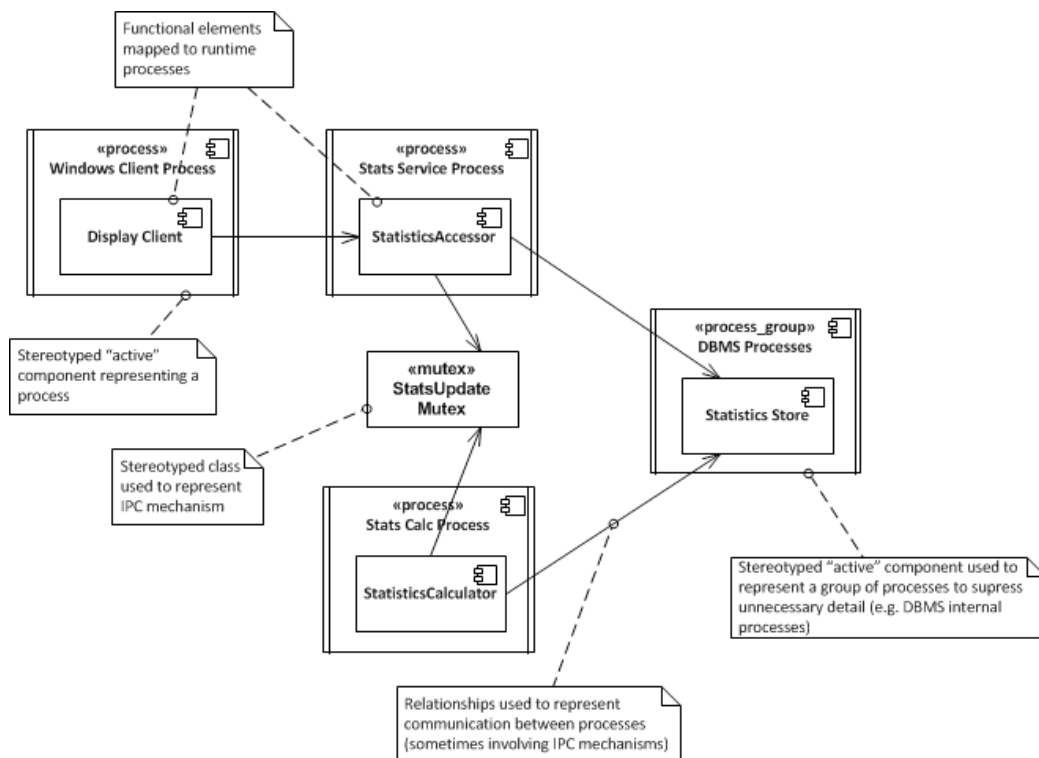


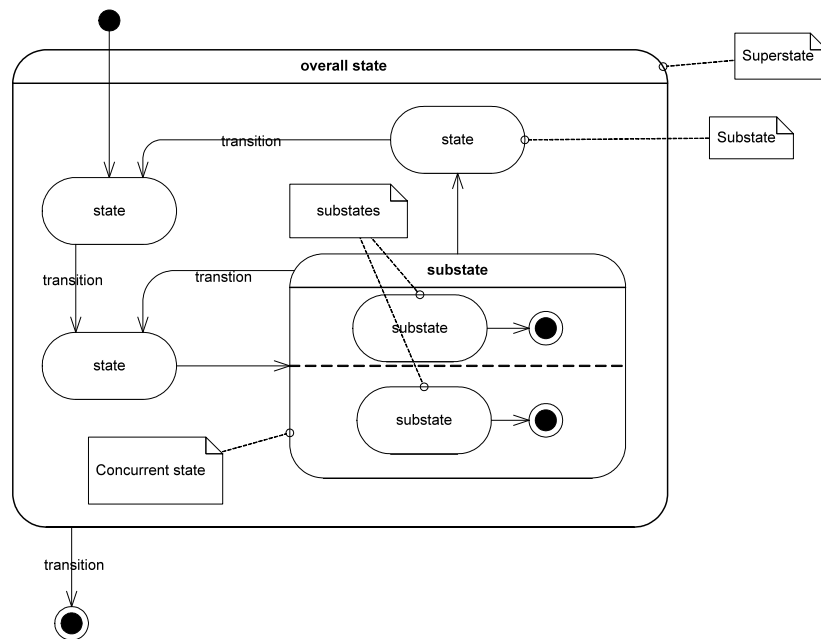
Figure 7. Concurrency Model

### 8.4.2 State Model

Model the states that the systems runtime elements can be in, the transitions between those states and the events which drive those transitions.

A state is an identified, named stable condition which occurs during the system's runtime. An event is something that happens which causes an element to undergo a transition from one state to another. Actions may also be associated with transitions, so that while the element changes state, the action is performed.

Focus on a few key elements whose states and transitions help to illuminate key features of the architecture.



**Figure 8. State Model**

## 8.5 Deployment View

The Deployment view of the system defines the important characteristics of the system's operational deployment environment. This view includes the details of the processing nodes that the system requires for its installation (i.e. its runtime platform), the software dependencies on each node (such as required libraries) and details of the underlying network that the system will require.

### 8.5.1 Runtime Platform Model

Show the system's runtime platform (defining nodes, links and the mapping of functional elements or processes to nodes).

You can use a UML deployment diagram here, or a simpler boxes-and-lines diagram.

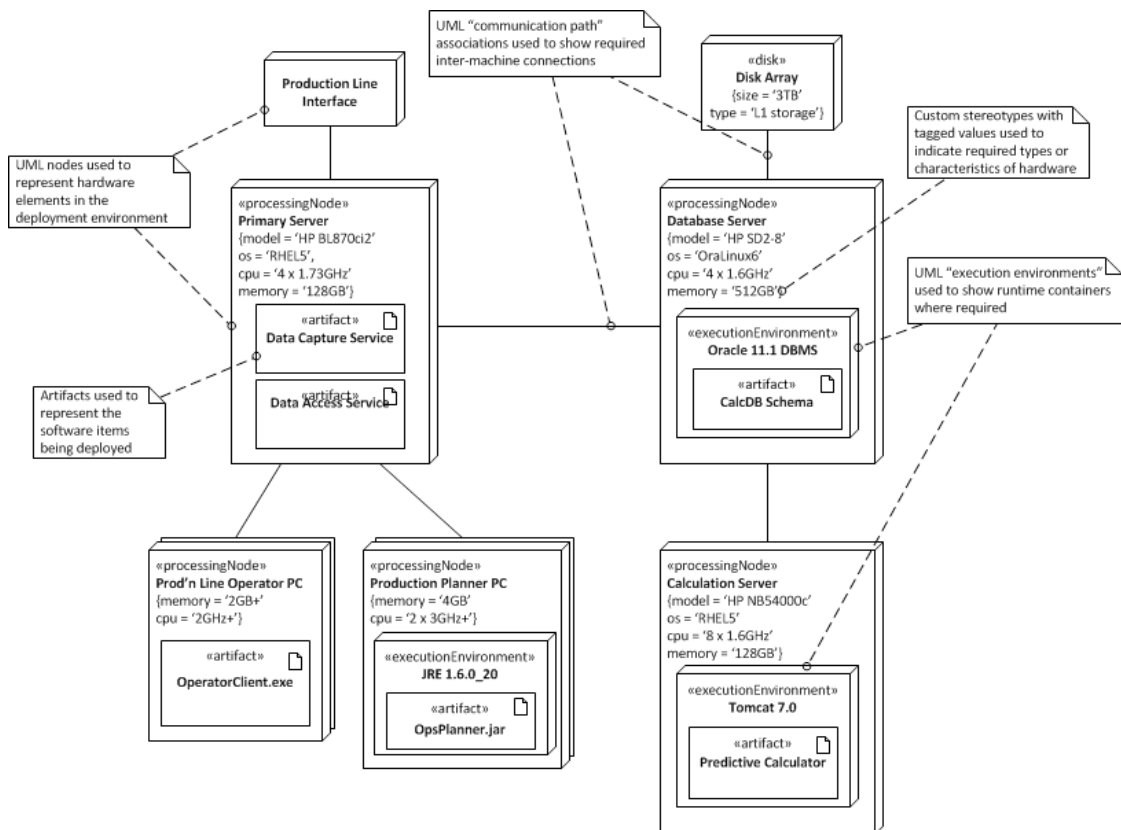


Figure 9. Deployment Model

It is often useful to explicitly map the functional elements onto the nodes that they will be running on, particularly if the deployment model is complex or the mappings aren't obvious.

Functional Element	Deployment Node(s)
element	node(s)
element	node(s)

### 8.5.2 Software Dependencies

Define the software that will be required on the various types of node in the runtime platform model, in order to support the system (such as operating system , system software or library requirements). Where versions are known you should state these.

Clearly state any known version dependencies (eg component A requires at least version X of component B).

This can usually be presented in tabular form.

### **8.5.3 Network Model**

If network requirements are complex, include a network model that illustrates the nodes, links and network hardware that the system requires, making quality of service requirements clear.

## 8.6 Development View

The Development view of the system defines any constraints on the software development process that are required by the architecture. This includes the system's module organisation, common processing that all modules must implement, any required standardisation of design, coding and testing and the organisation of the system's codeline.

*Much of the information in this view is normally presented at a summary level, with more detail being available in other developer focused documents such as a development standards document. However you may still need to record some architecturally significant decisions at this stage, for example around choice of libraries or frameworks, or approach and tools for software deployment or configuration management.*

### 8.6.1 Module Structure

*Use a model that defines the code modules that will be created and the dependencies between them. A UML package diagram is often an effective way to achieve this.*

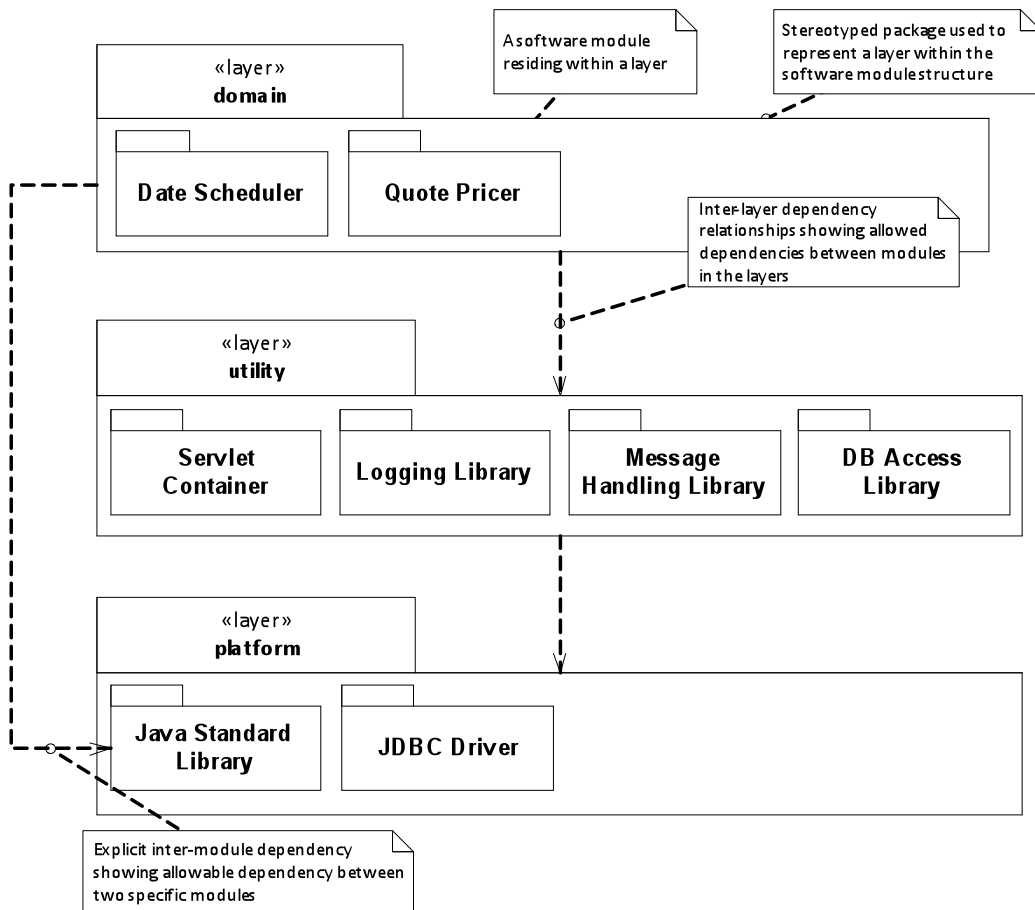


Figure 10. Module Structure Diagram

### 8.6.2 Common Design

*Define the common design (such as logging, security, tracing and so on) that must be performed in a standard way across the system and how it should be performed (e.g. via a design pattern or reference to a code library or sample).*

### **8.6.3 Standards for Design, Code and Test**

Define any standards that must be followed for design, code and unit testing, probably by reference to an external document.

### **8.6.4 Codeline Organisation**

Define the codeline structure (i.e. how the source code will be held as a directory hierarchy and how it will be built into deliverable software). Define the directory hierarchy, build tools and delivery tools (such as testing or continual integration tools) that will be used to deliver the software for testing and production.

## **8.7 Operational View**

The Operational view defines how the system will be installed into its production environment, how data and users will be migrated to it and how it will be configured, managed, monitored, controlled and supported once this is achieved. The aim of the information in this view is to show how the operational environment is to be created and maintained, rather than to define detailed instructions or procedures.

### **8.7.1 Installation and Migration**

Define the high-level steps required to install the system and any specific or unusual requirements for it.

If parallel running of old and new systems is required, explain how this will be done without disrupting existing systems, and the transition states required.

### **8.7.2 Operational Configuration Management**

Define the main groups of operational configuration items and common sets of values for them (e.g. batch and overnight sets) and explain how these groups will be managed in the production environment.

### **8.7.3 System Administration**

Explain the requirements the system places on the systems administrators (in both routine and exceptional situations) and the facilities that the system will provide or rely on in the operational environment.

### **8.7.4 Provision of Support**

Define the groups involved in providing support for the system and the roles and responsibilities of each (including escalation procedures if relevant).

## 9. System Qualities

This section explains how the architecture presented meets its each of its required system quality properties.

*While much of this information will be intrinsic to the views documented in the previous chapter, it is often useful to bring out some of it separately. In particular, if a quality property such as security or performance depends on features documented in several different views, then you should explain this here. For example, scalability may depend on optimisations in the data model (documented in the Information View) along with load balancing components (documented in the Deployment View).*

### 9.1 Performance and Scalability

*For each of the main performance and scalability requirements, explain how the system will meet the requirement. Refer to practical testing and performance modelling work that has been performed as part of applying this perspective.*

Requirement	How Met
1. average user response time should be XX under load YY	refer to performance modelling spreadsheet

### 9.2 Security

*For each of the main, security requirements, explain how the system will meet the requirement. Define (or reference) the threat model, security policy and security design that have been used as part of applying this perspective.*

Requirement	How Met
1. all users must be authenticated before being allowed to access the system	access to all screens is via standard login screen with passwords synchronised overnight to central LDAP service

### 9.3 Availability and Resilience

*Explain the A&R requirements.*

*Define the availability schedule(s) for the system.*

*Explain how the system will meet the requirements, referring to practical testing, modelling and design work that has been performed as part of applying this perspective.*

Requirement	How Met
1. There should be no single point of failure	all deployment nodes are clustered or load-balanced; where nodes are clustered, component failure is detected automatically and the passive node is brought up automatically

### 9.4 Evolution

*Explain the evolution requirements.*

*Define the evolutionary dimensions that are relevant to the system.*

Explain how the system will meet the requirements, taking into account the likelihood of each type of evolution occurring (explaining how the probabilities were arrived at) and referring to the design work performed as part of applying his perspective.

Requirement	How Met
1. it must be possible to add extra input channels without having to redesign the core system	input channel components are loosely coupled to central processing modules via standardises abstract interface

## **9.5 Other Qualities**

### **9.5.1 Accessibility**

Explain how the system meets any accessibility requirements (if any).

### **9.5.2 Internationalisation**

Explain how the system meets any internationalisation (or localisation) requirements (if any).

### **9.5.3 Location**

Explain how the system meets any requirements for the geographical location(s) it is to be installed in (if any).

### **9.5.4 Regulation**

Explain how the system meets any regulatory requirements (if any).

### **9.5.5 Usability**

Explain how the system meets any usability requirements (if any).

## 10. Appendices

### 10.1 Appendix: Decisions and Alternatives

Explain the primary architectural design decisions that have been made, their rationale and the set of alternatives considered for each.

### 10.2 Appendix: Questions and Answers

As you develop your architecture and explain it to people, the same questions are probably going to be asked repeatedly. In some cases, this is a signal to improve one particular part of the architectural description, but in others there is no one place that can address the query, so answer these questions here.

It is also sometimes the case that there are questions that you want readers to ask themselves in order to ensure that they have understood the architecture and its implications and a good way of encouraging this is also to put an appropriate question and answer in this section.

### 10.3 Appendix: References

List the references that you use throughout the text to allow them to be cited and located. For example:

- [1]. Software Systems Architecture: Working with Stakeholders Using Viewpoints and Perspectives, 2<sup>nd</sup> Edition. Nick Rozanski and Eoin Woods, Addison Wesley, 2011.