Annex to Vetronics Standards & Guidelines: VSI Metrics for Electronic Architecture Assessment

S J Press
QINETIQ/TS/FPPS/TR0900176
August 2009

The investigation which is the subject of this Report was initiated by the FRES IPT, DE & S, MOD Abbey Wood, Walnut Floor 3, Bristol BS34 8JH and was carried out under the terms of Dstl Purchase Order DSTLX-1000007039 Amendment 2 dated 12/11/2008.

Administration page

Customer Information	
Customer reference number	
Project title	FRES SV
Customer Organisation	FRES IPT
Customer contact	M Kellaway
Contract number	Dstl Purchase Order DSTLX-1000007039 Amendment 2
Milestone number	N/A
Date due	August 2009

Principal author	
S J Press	S J Press
Room G081, Bldg A5, Cody Technology Park, Farnborough, Hants, GU14 0LX	Sjpress1@QinetiQ.com

Release Authority	
Name	M Lemon
Post	Project Manager QinetiQ
Date of issue	February 2009
Name	M Uttley
Post	Project Manager DSTL
Date of issue	August 2009

Record of changes		
Issue	Date	Detail of Changes
Issue 1	Aug 2009	Initial Issue

List of contents

1 Introduction

1.1 Aim

The aim of this document is to outline a set of metrics to allow an assessment of an Electronic Architecture built using VSI (Vehicle System Integration) Standards and Guidelines [1]. In addition to defining the metrics some guidance is given on possible scoring systems to aid assessment, and some examples of the detailed parameters for assessment of highly complex platforms is provided in an Appendix.

1.2 General

The VSI metrics are intended to help achieve a relatively objective assessment of the extent to which a vetronics architecture is "VSI compliant". The metrics provide a mechanism for assessing compliance (based on considering a number of relevant aspects with a scoring mechanism for each) and deriving an overall score from the individual aspect scores.

The metrics have been derived by and are endorsed by members of the VSI Steering Group, which includes representation by QinetiQ, BAE Systems, General Dynamics (UK), Selex SAS, Ultra Electronics, Thales Air Defence and Thales Land & Joint.

Each assessment will be different since it will be based on particular requirements, scenarios and assumptions. These aspects need to be defined for each assessment, which can be seen as a tailoring of the metrics. Each score is dependent on the tailoring performed.

Hence statements like "Architecture A is VSI compliant" or "Architecture A scores 3.6" are – without further qualification defining the tailoring – meaningless.

While they provide a degree of objectivity, the metrics still require judgement to be applied and so are not truly objective.

> These factors mean that the metrics need to be evaluated by individuals with the necessary technical understanding of the context (requirements, etc) and of the architecture being assessed. Assessors also require an appreciation of the assessment process itself and what the results do (not) mean.
>
> Without this foundation, assessment will be a "garbage in, garbage out" activity.

These statements also imply that – if it is to be useful – an assessment is not a trivial undertaking. It requires planning, forethought and appropriate effort.

Based on what the VSI Standards and Guidelines are intended to achieve, a number of different architecture characteristics have been identified for assessment. These characteristics are:

- **Reconfigurability** – The extent to which the architecture allows changes to the system in the field, e.g. the substitution of one system element with another of similar performance (Adaptability); the interchange of system elements between different types of platform (Interchangeability).

- **Enhanceability** – The extent to which the architecture facilitates major changes to the larger system, e.g. fitting a new weapon system, new sensor type or additional new subsystem. This is subdivided into:
  - Capacity - This sub-dimension addresses the spare capacity that exists for key aspects associated with system enhancement (e.g. physical space, power, cooling, bandwidth or operator workload);
  - Modularity - The extent that the modularity of the architecture allows the system to be upgraded;
  - Enablers - The availability of skills, knowledge, tools etc. needed to enhance the architecture;
- **Integration** – The extent to which the EA provides the ability for data to be received and transmitted to other electronic systems, and to allow control of the EA.:
  - Internal to the platform (to other EA components);
  - External to the architecture under review (any direct interface to other systems e.g. Network Enabled Capability (NEC));
  - System Control – ensuring that users (human or automated) can control all relevant platform resources via the architecture to achieve their objectives.
- **Logistics Support** – The support provided by the architecture to Logistics activities;
  - Built In Test (BIT);
  - Integrated Logistics Support (ILS) Data Transfer.
- **Scalability** – the measure with which the performance of the system can increase or decrease in response to changes in resource demand:
  - Vertical scalability – additional resources added to existing system elements e.g. more memory, processor upgrade.
  - Horizontal scalability – the ability to add to or remove system elements to scale the system to meet differing performance requirements of different variants;
- **Openness** – An assessment of the system to ensure that all relevant aspects are open. This assessment includes:
  - Standards & Technology selection – To ensure that the VSI recommended technologies are used where possible, or open standards are used as an alternative.
  - Documentation – for hosting or interacting with other systems or processes.
  - Interface Control Document (ICD) – Fully populated and completed Interface Control Documents for all software/hardware interfaces.

Together, these characteristics may be seen as representing the stretch potential of the architecture.

For each characteristic, a score from 1 to 5 is derived and these scores can then combined into other scores. Most characteristics are decomposed into more detailed elements which are also scored.

This document describes the following:

- How the VSI metrics can be used;
- The metrics in more detail;
- How individual metrics scores can be combined.

2 Using the metrics

There are a number of key points to be considered when using the metrics.

2.1.1 Defined context and rationale

Any use of the metrics should involve a clear statement of:

- the scope of the assessment – i.e. the boundary of the architecture/system;
- the assumptions associated with each metric (see details below);
- the time for which the assessment is made (eg as now or as in 2012);
- the assumptions on which a future (eg as in 2012) assessment are based, such as that between now and then weapon X has been added to the system;
- the key facts/assumptions that justify each assigned score (unless entirely obvious) – including the impact of uncertainty;
- the weightings applied to each score, and the rationale for these.

These aspects are discussed below.

2.1.2 Assessment subject

The metrics may be used to assess a proposed architecture and system, or an existing one. Obviously, an existing architecture is much more tangible and so the assessment can be carried out with more certainty.

Although some fairly obvious "translation" may be required in some places, the metrics can also be used to assess parts of an architecture. They may also be used to assess the compliance of items that interface with the architecture, eg a "VSI compliant gun", although for this rather more modification of the metrics may be required. This extended use has not so far been fully explored so the details of it are undefined.

2.1.3 Context

A crucial fact to consider is that the metrics are largely context-dependent. That is, they assess compliance for an assumed set of conditions. For example, Reconfigurability is assessed against an assumed set of system elements to be interchanged. If a different set is assumed, the result may well be different. Thus a statement such as "X has a VSI compliance score of N" is only meaningful if the basis of the assessment is also stated. Further, the top level score is obtained from more detailed scores by applying weightings, and different weightings will clearly give different end results.

2.1.4 Time point

Some of the characteristics assessed are likely to change over time. Spare capacity may be used up as the system changes, the demands for connectivity with interoperating systems may grow, what was a widely-supported open standard may become obsolete, and what was a readily-available skill may become scarce. All these factors (spare capacity, etc) are reflected in the metrics.

This means that as well as assessing an architecture as it is now (or is planned to be shortly), it is desirable to assess the situation as it will be in N years' time. Such assessment cannot, of course, be precise but as long as it is based on clear, explicit assumptions, it can offer a valuable insight.

The result may be a profile such as that below.

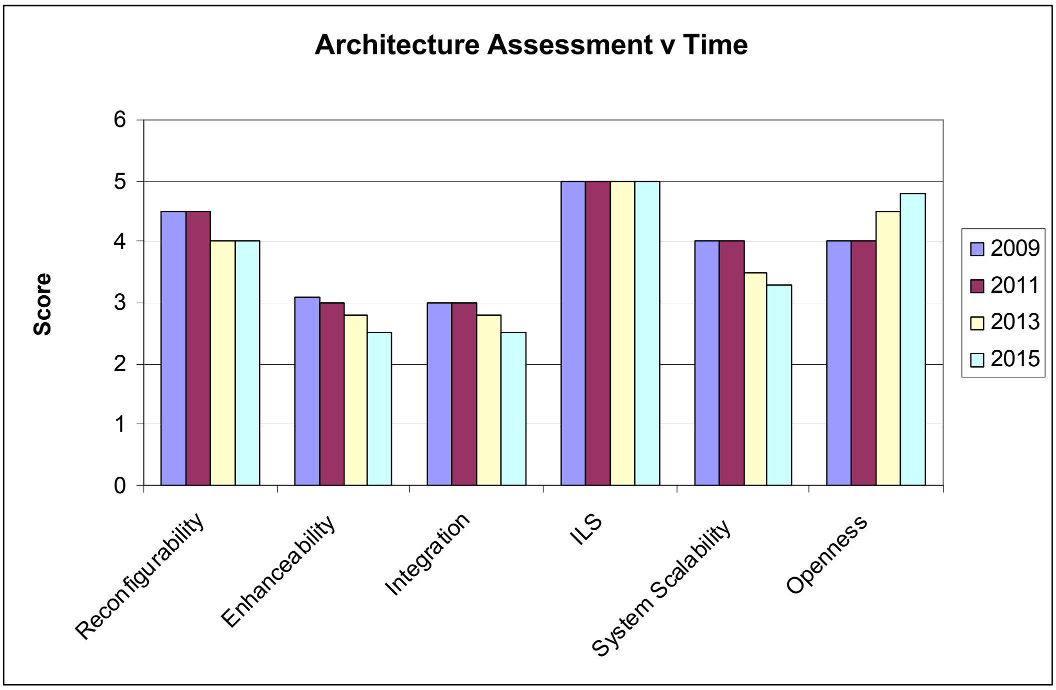

Figure 1 – Typical Architecture Assessment versus Time

This profile is, of course, just for illustration. In general, one might expect scores to decrease over time, but some may increase, e.g. as a particular standard becomes more adopted or as processor and communications technology advances.

A profile of just the overall score for two different architectures may also be informative. Architectures A and B may the same assessed as of now, but significantly different over time.

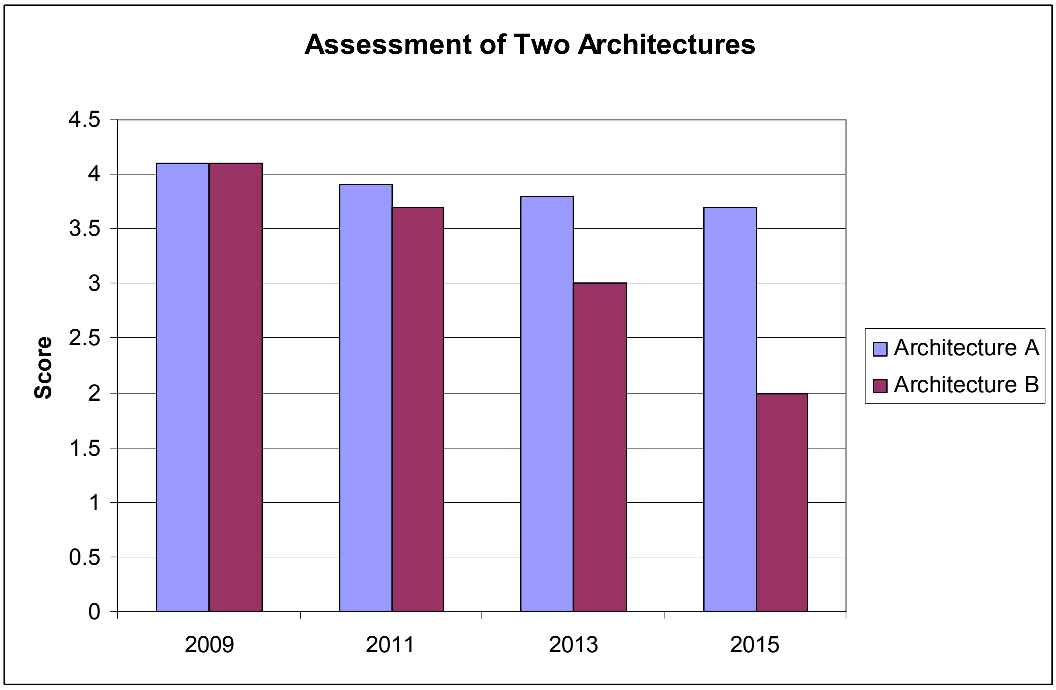

Figure 2 – Typical Assessment of Two Architectures

2.1.5 Confidence level

Often, an assessment will need to be made using information that is uncertain. This may relate to the architecture being assessed (e.g. it does not yet exist, or it exists but information about it is incomplete) and/or to the assessment criteria (e.g. the growth margin that is expected to be needed). Scores may be adjusted to reflect any uncertainty, but the way this is done and the rationale needs to be clear.

2.1.6 Other factors

It should be clear from the brief introduction to them above that the metrics offer a characterisation, from one perspective, of a system and its electronic architecture. Clearly, there are many other perspectives: cost, performance, reliability, robustness, etc. Thus for most decisions, the metrics will be just part of a bigger picture. The weight attached to the metrics may vary in different circumstances and will depend on priorities, e.g. the relative importance of key requirements defined by a customer.

2.1.7 Over-simplification

At their most stark, the metrics provide a single number representing VSI compliance, but this is often of limited use. For example, if architecture A scores 4.2 and B scores 3.5, this reflects that A is more compliant. However, it does not – in itself – tell one why A is more compliant or suggest how difficult it would be to enhance B's compliance.

The single number needs to be used with care and it is often better to consider the individual metrics (for Reconfigurability, etc) and their constituent parts. Profiles of these individual scores can be formed as discussed in section 4. It is these detailed scores that will be most useful in dialogue about compliance. Setting a "pass mark" for compliance is possible, but simplistic. It may be appropriate to do so when comparing alternatives against a fixed set of assumptions, although even here a relative approach (comparing the score for each alternative) may be better than an absolute one of setting a threshold.

3 Metrics in detail

3.1 Metrics overview

The high level metrics outlined above can be seen as "dimensions" of the highest level metric: VSI Compliance. Each is in turn composed of some lower level sub-dimensions which are assigned a score. The score (on a 1 to 5 scale) is assigned by considering two lowest level "parameters".

This is shown in the following generalised diagram.

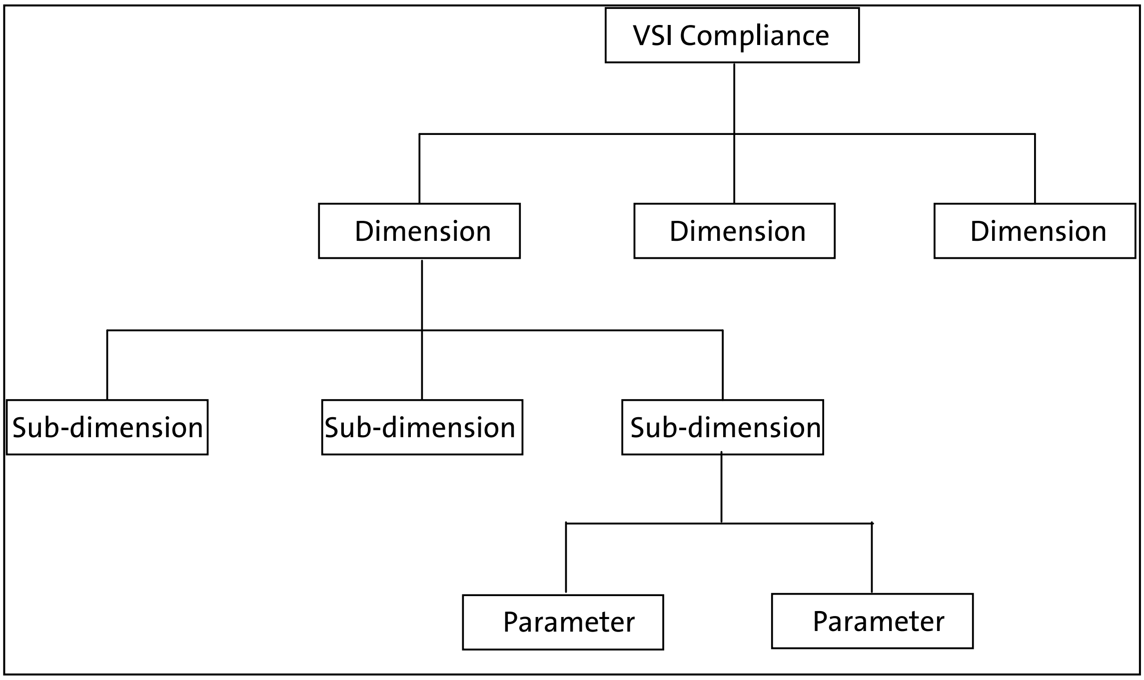

Figure 3 – Metrics Dimension Hierarchy

The actual metrics/dimensions are shown in the following table (rather than in tree form for convenience).

	Dimension	Sub-dimension
VSI Compliance	Reconfigurability	Adaptability
		Interchangeability
	Enhanceability	Capacity
		Modularity
		Enablers
	Integration	Internal Platform Data Provision
		External Platform Data Provision
		System Control
	Integrated Logistics Support	BIT
		ILS Data Transfer
	System Scalability	Vertical scalability
		Horizontal scalability
	Openness	Standards & Technology selection
		Documentation
		ICDs

Table 1 – VSI Compliance breakdown

The score for the lowest level dimension is derived from a table that indicates the situation (in terms of the values the parameters take) corresponding to each score from 1 to 5. This is shown below.

5	4	3	2	1
Some narrative describing 'Par 1 Value 1' for Parameter 1 and 'Par 2 Value 1' for Parameter 2.	Some narrative describing 'Par 1 Value 2' for Parameter 1 and 'Par 2 Value 1' for Parameter 2 OR 'Par 1 Value 1' for Parameter 1 and 'Par 2 Value 2' for Parameter 2.			

Table 2 – Generic dimension parameter table type 1

An alternative representation provides a score for each possible combination of values for the two parameters related to the dimension, as shown below. (Note that the number of distinct values for each parameter varies between dimensions and is not always the same for each parameter. Also, the allocation of the 1 to 5 score to cells obviously varies.)

	Parameter 1		
Parameter 2	'Par 1 Value 1'	'Par 1 Value 2'	'Par 1 Value 3'
'Par 2 Value 1'	5	4	2
'Par 2 Value 2'	4	3	2
'Par 2 Value 3'	1	2	1

Table 3 – Generic dimension parameter table type 2

Either version can be used for scoring, but the first version provides fuller definitions of the scores. The second version may be used when the assessor is very familiar with the definitions and/or may be used to actually record scores by marking the relevant cell.

3.2 Description of the Metrics

3.2.1 Reconfigurability

Reconfiguration relates to changing the platform to meet short-term needs. It is typically characterised by being performed in the field, over a short timeframe, using the available skills and tools, and being temporary, i.e. easily reversible. Larger scale changes are addressed under Enhanceability.

Adaptability

The first dimension of Reconfigurability is Adaptability, which focuses on changing the system in an essentially like-for-like way, e.g. to replace sensor A by sensor B.

The parameters for Adaptability are

- Scope: the proportion of relevant changes that can be made in the field (i.e. with available skills and tools, in an acceptable time);
- Difficulty: the work needed to make the changes, in terms of modification to various parts of the system/platform.

It will typically be necessary to define

- the (kind of) changes that are relevant for a particular assessment;
- the (kind of) modifications that are considered minor.

The following table defines the score.

5	4	3	2	1
All relevant changes can be made in the field (i.e. with available tools and skills and in an acceptable time) without making any modifications to parts of the system/platform.	All relevant changes can be made in the field with only minor modifications needed. OR The majority of relevant changes can be made in the field without making any modifications.	The majority of relevant changes can be made in the field with only minor modifications needed.	The minority of relevant changes can be made in the field with no, or only minor, modifications needed.	No relevant changes are possible in the field.

Table 4 – Adaptability Scoring Matrix

Note that it is a common characteristic of these metrics that the best aspect is scored. For example, suppose there are six relevant changes and four can be made without any modification, while the remaining two require minor modifications. This would score 4 (the majority can be done with no modification) rather than 2 (the minority can be done with only minor change).

It is also common that every possible situation is not covered in detail. In the above example, the score would be 4 whether the remaining two changes can be made with minor modifications or are impossible.

The metrics essentially reflect the "highlights" of the situation.

Alternatively, the following table shows how the scores are allocated in relation to parameter values.

	Scope			
Difficulty	All	Majority	Minority	None
No modification	5	4	2	
Minor Modifications	4	3	2	
Not Possible in Field				1

Table 5 – Alternative Adaptability Scoring Matrix

Interchangeability

Interchangeability is about the number of relevant system elements ("boxes" or LRUs) that can be moved between relevant platform types to perform essentially the same function on the new platform.

The parameters for Interchangeability are

- Scope: the proportion of relevant system elements that can be moved;
- Difficulty: the work needed to make the move, in terms of modification to parts of the system/platform.

It will typically be necessary to define

- the (kind of) system elements that are relevant for a particular assessment;
- the (kind of) modifications that are considered minor;
- the relevant platform types.

The following table defines the score.

5	4	3	2	1
All relevant system elements can be moved between all relevant platform types in the field (ie with available tools and skills and in an acceptable time) without making any modifications to the system/ platform.	All relevant system elements can be moved between all relevant platform types in the field with only minor modifications needed to the system/ platform. OR The majority of relevant system elements can be moved between all relevant platform types in the field without making any modifications to the system/ platform.	The majority of relevant system elements can be moved between all relevant platform types in the field with only minor modifications needed to the system/ platform.	The minority of relevant system elements can be moved between all relevant platform types in the field with no, or only minor, modifications needed to the system/ platform.	No relevant system elements can be moved in the field.

Table 6 - Interchangeability Scoring Matrix

Alternatively, the following table shows how the scores are allocated in relation to parameter values.

	Scope			
Difficulty	All	Majority	Minority	None
No modification	5	4	2	
Minor Modifications	4	3	2	
Not Possible in Field				1

Table 7 –Alternative Interchangeability Scoring Matrix

3.2.2 Enhanceability

Enhanceability relates to major system changes, typically made by the original supplier (or equivalent), in specialised workshops while the platform is out of use for a relatively long period.

Capacity

The Capacity dimension addresses the spare capacity that exists for key aspects associated with system enhancement. Relevant aspects may include physical space, power, cooling, or bandwidth. Capacity is assessed for the current architecture/system; potential for easily extending it is assessed under Scaleability. Capacity is compared with that expected to be needed over a defined period.

The parameters here are

- Scope: the proportion of relevant capacity aspects (space, etc) for which there is spare capacity;
- Margin: the spare capacity that exists, compared with an estimated need based on experience and projected future changes.

Note that the key concern is having enough spare capacity (growth margin) for future needs. This may obviously imply over-supply initially, which in turn may imply additional initial cost. The expected benefit is reduced cost later, when enhancement is needed, and/or increased capability (since not having enough spare capacity may actually make some enhancement infeasible).

It will typically be necessary to define

- the relevant capacity aspects;
- the predicted spare capacity needed for each aspect;
- the basis for the predicted need, including the timescale;

The following tables show how the scores are allocated.

5	4	3	2	1
For all relevant aspects (interface bandwidth, processing, protocol /message structure capacity, power supply, physical space, etc) of the system, the architecture has more spare capacity than is needed, as suggested by experience and predictions of future need.	For all relevant aspects (interface bandwidth, etc) of the system, the architecture has enough spare capacity, as suggested by experience and predictions of future need. OR For the majority of relevant aspects (interface bandwidth, etc) of the system, the architecture has more spare capacity than is needed, as suggested by experience and predictions of future need.	For the majority of relevant aspects (interface bandwidth, etc) of the system, the architecture has enough spare capacity, as suggested by experience and predictions of future need.	For only the minority of relevant aspects (interface bandwidth, etc) of the system, does the architecture have enough spare capacity, as suggested by experience and predictions of future need.	The architecture does not have enough spare capacity for any of the relevant aspects.

Table 8 – Capacity Scoring Matrix

	Scope	
Margin	All	Majority
Excess	5	4
Enough	4	3
Too Little	1	2

Table 9 – Alternative Capacity Scoring Matrix

Modularity

The modularity sub-dimension addresses the way in which the elements of the architecture are designed to allow further upgrades and additions. For different parts of the architecture modularity may be defined at different levels e.g. LRU level for many system elements, or card level for rack mounted equipment.

Software modularity should also be considered, this being of great relevance when middleware is used to abstract the higher and lower levels of software.

The parameters are:

- The extent to which the architecture elements are modular.

It will typically be necessary to define:

- The granularity of modularity expected(big “chunks” or small, LRU or software element)
- Any particular way in which modularity is assessed.

5	4	3	2	1
The architecture is highly modular	The architecture is largely modular but small parts are not.	Parts of the architecture are modular, but parts are not.	Only small parts of the architecture are modular	The architecture is not modular

Table 10; Modularity Scoring Matrix

Modularity				
Highly	Large part	Around half	Small part	None
5	4	3	2	1

Table 11; Alternative Modularity Scoring Matrix

Enablers

The Enablers dimension assesses the availability of the skills, tools, knowledge, etc needed to enhance the architecture.

The parameters are:

- Skills, etc : the extent of skills, knowledge, tools, etc needed to change the architecture;
- Availability : the extent to which the skills etc exist. Skills etc may be available in depth, in a limited way (e.g. just one small team has them), or not at all.

Note that the ease (or otherwise) with which missing skills etc. might be obtained is not reflected.

It will typically be necessary to define

- the relevant skills, etc and those that are considered important;
- the way in which availability is judged, e.g. what constitutes availability in depth or limited availability;
- the scope of availability, e.g. globally, within the UK, within the organisation(s) expected to perform enhancements.

The following tables show how the scores are allocated in relation to parameter values.

5	4	3	2	1
All relevant skills, etc are available in depth.	All the important relevant skills, etc are available in depth.	Some important relevant skills, etc are available in depth. OR All important relevant skills, etc are available but only in a limited way.	Some important relevant skills, etc are available but only in a limited way.	None of the relevant skills, etc are available.

Table 12 – Skills Scoring Matrix

	Skills, etc		
Availability	All	All Important	Some Important
In Depth	5	4	3
Limited	3	3	2
Unavailable	1	1	1

Table 13 – Alternative Skills Scoring Matrix

3.2.3 Platform Integration

The extent to which the EA provides the ability for data to be distributed and used between the architecture under review and other electronic systems. This relates to data provision internally, externally and to allow control of the system.

Internal data provision

This sub-dimension relates to the data provision (both sending and receiving data) between the architecture under review and the rest of the electronic platform architecture elements. The data should be transmitted or received in a secure and timely manner.

The parameters are:

- The ability to transmit information to other logical architectures on the platform (this may not be a direct link but via another architecture such as a backbone)
- The ability to receive information from other logical architectures on the platform (this may not be a direct link but via another architecture such as a backbone).

It will typically be necessary to define:

- The relevant types of data.
- Any relevant criteria of the data for security, safety criticality, timeliness etc.

5	4	3	2	1
All relevant information is available in a timely and secure manner	The majority of the data is transmitted and received in a timely and secure manner	The majority of the data is transmitted but not received in a secure and timely manner OR The majority of the data is received but not transmitted in a secure and timely manner	Only a minority of the data is transmitted or received in a secure and timely manner	The architecture is not capable of transmitting data in a timely and secure manner

Table 14; Internal Data Provision Scoring Matrix

	Data Transmitted			
Data Received	All	Majority	Minority	None
All	4	3	2	1
Majority	3	3	2	1
Minority	2	2	2	1
None	1	1	1	1

Table 15; Alternative Internal Data Provision Scoring Matrix

External platform data provision

This sub-dimension relates to the data provision (both sending and receiving data) between the architecture under review and electronic systems external to the electronic architecture under review. This would include wider off-platform NEC applications. This may be accomplished via a suitable communications gateway to existing communications equipment or systems, or directly to new communications systems. The data should be transmitted or received in a secure and timely manner.

The parameters are:

- The ability to transmit information to electronic equipment external to the EA under review e.g. communications gateway, or directly to any other NEC device fitted on the platform (this may not be a direct link but via another architecture such as a backbone)
- The ability to receive information from electronic equipment to the EA under review e.g. communications gateway, or directly from any other NEC device fitted on the platform (this may not be a direct link but via another architecture such as a backbone).

It will typically be necessary to define:

- The relevant types of data.
- Any relevant criteria of the data for security, safety criticality, timeliness etc.

5	4	3	2	1
All relevant information is available in a timely and secure manner	The vast majority of the data is transmitted and received in a timely and secure manner	The majority of the data is transmitted but only a minority of data received in a secure and timely manner OR The majority of the data is received but minority of data transmitted, in a secure and timely manner	Only a minority of the data is transmitted or received in a secure and timely manner	The architecture is not capable of transmitting data in a timely and secure manner

Table 16; External Data Provision Scoring Matrix

	Data Transmitted			
Data Received	All	Majority	Minority	None
All	5	-	-	-
Majority	-	4	3	-
Minority	-	3	2	-
None	-	-	-	1

Table 17; Alternative External Data Provision Scoring Matrix

System control

This dimension addresses how users (human or automated) can control relevant platform resources via the architecture to achieve their objective.

The parameter is

- Scope: the amount of relevant resources that can be controlled in a secure and safe manner under relevant conditions.

It will typically be necessary to define

- the relevant (types of) resources;
- criteria for security and safety, and any other relevant characteristic such as timeliness;
- the conditions of interest;

The following tables show how the scores are allocated in relation to parameter values.

5	4	3	2	1
All relevant resources can be controlled by users/subsystems in a secure and safe manner and with an acceptable quality of service under all relevant conditions.	The majority of relevant resources can be controlled by users/ subsystems in a secure and safe manner and with an acceptable quality of service under all relevant conditions.	Approximately 50% of relevant resources can be controlled by users/subsystems in a secure and safe manner and with an acceptable quality of service under all relevant conditions.	The minority of relevant resources can be controlled by users/ subsystems in a secure and safe manner and with an acceptable quality of service under all relevant conditions.	No relevant resources can be controlled by users/subsystems in a secure and safe manner and with an acceptable quality of service under all relevant conditions.

Table 18 – System Control Scoring Matrix

Scope				
All	Majority	50%	Minority	None
5	4	3	2	1

Table 19 – Alternative System Control Matrix

3.2.4 Integrated Logistics Support (ILS)

ILS addresses how the architecture provides the features needed for logistics support.

BIT

This sub-dimension is about support for BIT (Built In Test) across the EA. The majority (if not all) the major EA LRUs (Line Replaceable Units) should be capable of generating and transmitting BIT status.

The parameters are:

- The number of LRUs that generate BIT.
- The amount of relevant BIT data routing that is supported by the architecture.

It will typically be necessary to define:

- The relevant types of BIT data.

5	4	3	2	1
All main LRUs generate BIT and the architecture supports the routing of the BIT data.	The majority of main LRUs generate BIT and the architecture supports the routing of the BIT data.	A minority of main LRUs generate BIT and the architecture supports the routing of the BIT data.	A minority of main LRUs generate BIT and the architecture only has limited support for the routing of the BIT data.	The architecture does not support the routing of any relevant BIT data

Table 20; BITS Scoring Matrix

	LRU BIT generation		
BIT data routing	All	Majority	Minority
Complete	5	4	3
Limited	2	2	2
None	1	1	1

Table 21; Alternative BIT Scoring Matrix

Data transfer

This dimension is about support for transferring logistics data on/off the platform.

The parameter is

- Scope: the amount of relevant data whose transfer is supported by the architecture.

It will typically be necessary to define

- the relevant (types of) data.

The following tables show how the scores are allocated in relation to parameter values.

5	4	3	2	1
The architecture supports transfer of all relevant data.		The architecture supports transfer of some relevant data.		The architecture does not support transfer of any relevant data.

Table 22 – Data Transfer Scoring Matrix

Scope		
All	Some	None
5	3	1

Table 23 – Alternative Data Transfer Scoring Matrix

Note that if parts of an architecture, or logical architectures, are assessed individually then this metric will not always be applicable. When it is not applicable no mark (0) should be assigned.

3.2.5 Scalability

This is the measure with which the performance of the system can increase or decrease in response to changes in resource demand. Extension of the architecture can be achieved by exploiting or enhancing existing elements or by the addition of further system elements.

Vertical scalability

Vertical Scalability is defined as additional resources, including software, being available or added to existing system elements to increase their performance. Vertical Scalability addresses how the existing architecture can be extended to provide additional performance (bandwidth, processing power, etc) by exploiting existing spare capacity, via simple replacement, or via minor modification.

The parameters are

- Scope: the amount of increased performance that can be gained
- Difficulty: the changes needed to achieve the gain

It will typically be necessary to define

- the kind of performance (bandwidth, etc) that is relevant
- criteria for “significant” performance increase
- the kind of changes that are considered minor

The following tables show how the scores are allocated in relation to parameter values.

5	4	3	2	1
Significant increased performance is possible through exploiting existing spare capacity or through “form and fit” module replacement.	Some increased performance is possible through exploiting existing spare capacity or through “form and fit” module replacement. OR Significant increased performance is possible through exploiting existing spare capacity or through “form and fit” module replacement, plus some software changes.	Some increased performance is possible through exploiting existing spare capacity or through “form and fit” module replacement, plus some software changes. OR Significant increased performance is possible through making some minor changes (not “form and fit” module replacement).	Some increased performance is possible through making some minor changes (not “form and fit” module replacement). OR Significant increased performance is possible through making major changes.	Some increased performance is possible through making major changes.

Table 24; Vertical Scalability Scoring Matrix

	Scope	
Difficulty	Significant	Some
None or Module Replacement	5	4
None or Module Rep, plus Software	4	3
Minor but not "form and fit"	3	2
Major	2	1

Table 25; Alternative Vertical Scalability Scoring Matrix

Horizontal scalability

Horizontal scalability is the ability of the system performance to be scaled by the addition or subtraction of system elements. The changes to the number of system elements should be able to be completed with minimal intervention from the user.

The parameters are

- Scope: the amount of increased performance that can be gained by the addition of elements;
- Difficulty: the changes needed to achieve the gain

It will typically be necessary to define

- the kind of performance (bandwidth, etc) that is relevant
- criteria for "significant" performance increase
- the kind of changes that are considered minor

The following tables show how the scores are allocated in relation to parameter values.

5	4	3	2	1
Significant increases in performance across a number of areas are possible by the addition of different system elements. Adding system elements requires minimal user intervention	Significant increases in performance across a number of areas are possible by the addition of different system elements. Adding system elements requires some user intervention.	Increases in performance across a number of areas are possible by the addition of different system elements. Adding system elements requires either minimal or some user intervention.	Some increases in performance for certain aspects are possible through the addition of a small number of extra elements. Adding system elements requires either minimal or some user intervention.	The architecture cannot readily accommodate additional elements to increase performance.

Table 26; Horizontal Scalability Scoring Matrix

	Additional performance		
Element addition	Significant for number aspects	Some for number of aspects	Some for certain aspects
Minimal user intervention	5	3	2
User intervention needed	4	3	2
None	1	1	1

Table 27; Alternative Horizontal Scalability Scoring Matrix

3.2.6 Openness

Open standards are the key enabler for the long term support of electronic architectures.

Standards and technology selection

Within the VSI Standards and Guidelines lists of recommended standards are given for different aspects of the overall architecture. The EA for new platforms should be made up using the recommended standards where possible, especially those designated for future systems. Where VSI standards are not used preference should be given to open standards. This applies to both software and hardware.

The parameters are:

- VSI recommended standard;
- Open standard.

It will typically be necessary to define

- An open standard;
- Allowable justification for use of non-open standards.

The following tables show how the scores are allocated in relation to parameter values.

5	4	3	2	1
All standards and technologies used are from the VSI recommended lists OR The majority of standards and technologies used are from the VSI recommended lists and the remaining standards and technologies are open	The vast majority of standards and technologies used are from the VSI recommended lists OR The vast majority of standards and technologies used are combined from the VSI recommended lists or open standards AND Any non-open standards used are fully justified	The majority of standards and technologies used are combined from the VSI recommended lists or open standards AND Any non-open standards used are fully justified		The minority of standards and technologies used are from the VSI recommended list OR the minority of standards used are not fully open OR There is no justification for the use of non-open standards

Table 28; VSI Standards and Technology Scoring Matrix

	VSI Technology Lists/Open standards		
Non-open standard	All	Vast Majority	Majority
None	5	-	-
Justified	-	4	3
Non-justified	-	1	1

Table 29; Alternative VSI Standards & Technology Scoring Matrix

Documentation

The Documentation dimension relates to the availability of documentation needed to enhance the architecture and/or system.

The parameters are

- Scope : the proportion of the required aspects of the architecture for which high quality documentation exists;
- Accessibility : the extent to which the documentation is available to any organisation likely to need it to enhance the architecture/system.

It will typically be necessary to define

- those aspects for which documentation is required (probably 'all' by default);
- what is considered "high quality" documentation;
- the way in which accessibility is assessed. Barriers may include the need to pay a fee or the retention of Intellectual Property. Whether the assessment is purely "as is", or is based on some assumed action such as paying a fee, must be clear.

The following tables show how the scores are allocated in relation to parameter values.

5	4	3	2	1
There is high quality documentation for all relevant aspects of the architecture and it is fully accessible to all relevant organisations.	There is high quality documentation for all relevant aspects of the architecture but it is only partially accessible to all relevant organisations.	There is high quality documentation for the majority of relevant aspects of the architecture and it is fully accessible to all relevant organisations. OR There is high quality documentation for all relevant aspects of the architecture but none of it accessible to "third party" organisations.	There is high quality documentation for only a minority of relevant aspects of the architecture but it is fully accessible to all relevant organisations. OR There is high quality documentation for the majority of relevant aspects of the architecture but it is either only partly accessible to all relevant organisations, or is not accessible at all.	There is high quality documentation for only a minority of relevant aspects of the architecture and it is either only partly accessible to all relevant organisations, or is not accessible at all. OR High quality documentation does not exist for any of the relevant aspects of the architecture.

Table 30- Documentation Scoring Matrix

	Scope			
Accessibility	All	Majority	Minority	None
Full	5	3	2	
Partial	4	2	1	
Not at All	3	2	1	
N/A				1

Table 31-Alternative Documentation Scoring Matrix

ICDs

As a consequence of the modularity aspects of the architecture there are defined interfaces between the logical architecture components. In order for the EA to be able to be augmented or upgraded at a later date fully populated ICDs must be available for release to third parties. These ICDs may be hardware only, software only, or most cases cover both hardware and software i.e. as many of the ISO layers as applicable.

Although ICDs on their own are useful there may need to be some system level functional information, or data model, available in order to ensure correct operation of the equipment. This is especially so where middleware is employed. This information should be included along with the ICD.

The parameters are:

- Scope - does the ICD address all the necessary areas?
- Quality – is the ICD fully populated and understandable?

It will typically be necessary to define

- Where ICDs should be produced.
- The aspects which the ICD should cover (e.g. full 7 layer ISO model by default).

The following tables show how the scores are allocated in relation to parameter values.

5	4	3	2	1
There is a fully populated ICD accessible to all relevant organisations	There is a fully populated ICD but it is only partially accessible to all relevant organisations	The ICD is partially populated but covers the majority of the relevant aspects and is accessible to all relevant organisations	The ICD is partially populated but covers a majority of the relevant aspects and is partially accessible to all relevant organisations OR There is a fully populated ICD but it is not accessible to all relevant organisation	The ICD is partially populated but covers a majority of the relevant aspects and is not accessible to all relevant organisations OR The ICD is partially populated but does not cover a majority of the relevant aspects and is not accessible to all relevant organisations OR There is no populated ICD

Table 32; ICDs Scoring Matrix

	Scope of ICD		
Accessibility	Fully	Partial	None
Full	5	3	1
Partial	4	2	1
Not at all	2	1	1

Table 33; Alternative ICDs Scoring Matrix

3.3 Parameter definition example

As described in the preceding sections, to enable an assessment to be performed the parameters used for assessing each sub-dimension need to be defined. These sub-dimensions will be different for different programmes on different platforms, and will vary according to the complexity of the EA, or part of EA, that is under assessment.

A guidance example is given in Annex A for a highly complex new platform design, which has been broken down for assessment purposes.

4 Use of metrics scores

4.1 Combining scores

The tables for each sub-dimension indicate a score for each lowest-level dimension. These 15 values can themselves provide an indication of VSI compliance, and of good and weak aspects of a particular architecture, e.g. if viewed as a histogram. For example:

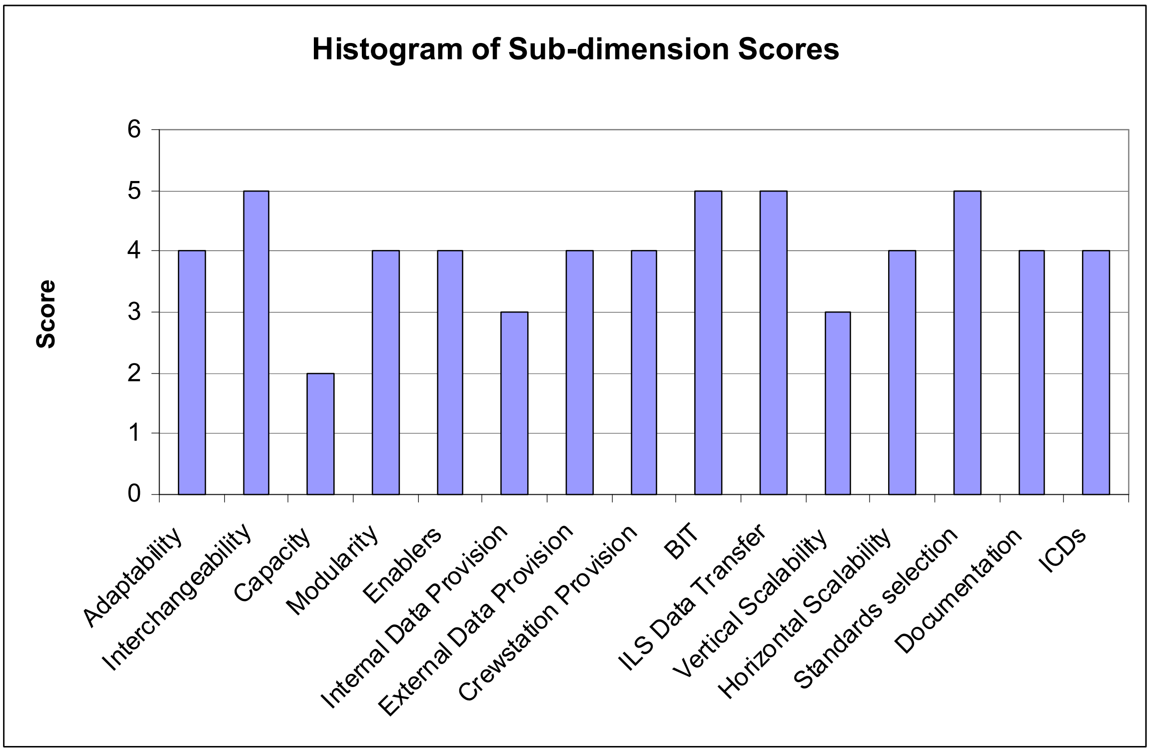

Figure 4 – Typical Histogram of Lowest Level Dimensions

However, a more abstract view based on the top-level dimensions may also be useful. Obtaining this clearly involves combining scores within each top-level dimension and this can be done via simple averaging or with a more sophisticated weighting system. A simple average of the above example gives:

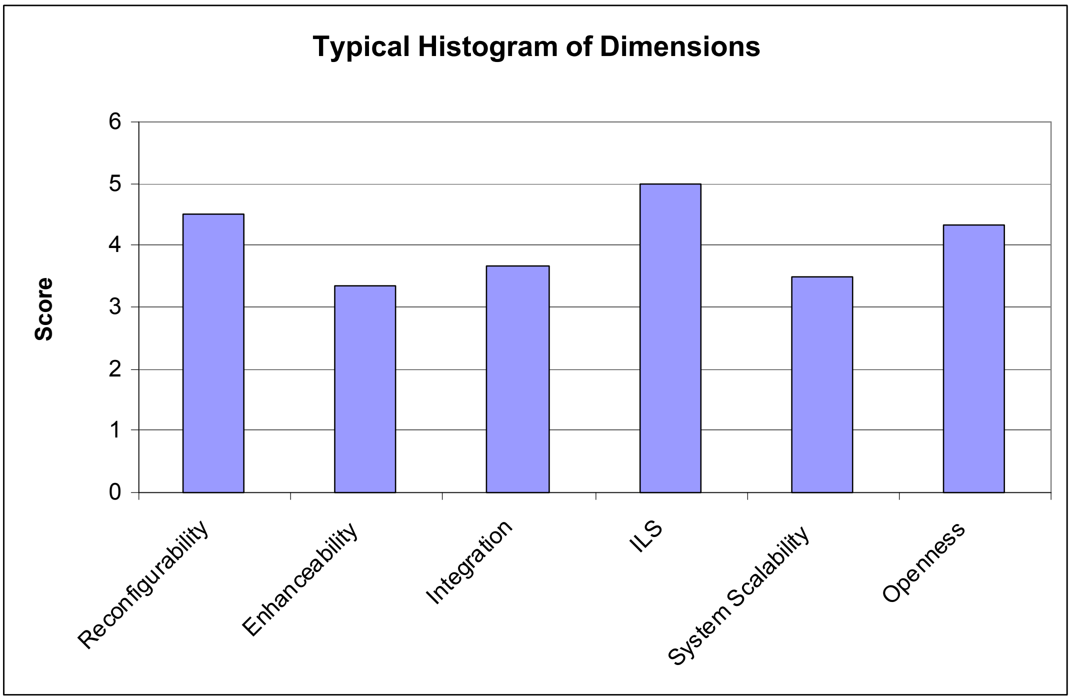

Figure 5 – Typical Histogram of Top Level Dimensions

Finally, a single score for VSI Compliance can be obtained from the top-level scores. This too can be a simple average or based on weightings. In the example, the simple average is 4. It should be noted that this score is dependent upon both the initial metrics assessment and the weightings. Consideration should always be given to undertaking a sensitivity analysis when using numerical scores to understand the likely error margin.

The weightings used obviously need to be decided upon, justified and defined for any assessment. A default weighting would be 0.16 for each top-level dimension, resulting in a simple average. However, it is better practice to actively consider the appropriate weightings for every particular assessment.

Weightings may be based directly on judgement using judgement panels or derived with techniques such as AHP (Analytic Hierarchy Process). An AHP approach compares each dimension with every other (e.g. Reconfigurability may be thought less important than Usability) and then derives the weighting appropriate for each dimension.

Where a number of different architectures solutions are being assessed using VSI metrics, as an alternative to simple averaging for score combination the Pugh concept selection technique can be used. See Annex B for a discussion on the application of these techniques to the multiple tender case.

The Openness dimension is central to VSI and the weighting attached to it may therefore be very high. Alternatively, a threshold may be set so that if, for example, a score of 4 or higher is not obtained for Openness, the architecture is considered to be non-compliant no matter what the other dimensions score.

4.2 Summary

Typically, VSI compliance will be one of several factors when considering an architecture or platform. VSI metrics provide a relatively objective way of assessing this compliance. Using the metrics is a non-trivial activity that should be undertaken with the necessary planning, skills and knowledge. The metrics are derived from six top-level dimensions that reflect the main drivers for VSI. Any assessment, and especially a comparison of alternative architectures, must be based on clear assumptions and criteria, since varying these may well result in different conclusions.

5 References

[1] Vetronics Standards & Guidelines, QINETIQ/EMEA/TS/CR0702540 Issue 3, August 2009.

6 Abbreviations/Acronyms

BIT	Built In Test
C2	Command and Control
DAS	Defensive Aids Suite
EA	Electronic Architecture
GFE	Government Furnished Equipment
HUMS	Health and Usage Monitoring System
ICD	Interface Control Document
ILS	Integrated Logistics Support
IPT	Integrated Project Team
LRU	Line Replaceable Unit
MilCAN	Military CANbus
NEC	Network Enabled Capablity
RAM	Random Access Memory
ROM	Read Only Memory
SRL	System Readiness Level
TRL	Technology Readiness Level
VSI	Vehicle Systems Integration

A Examples of definition of assessment parameters

A.1 Worked example for a highly complex EA for a new platform design

The VSI Standards and Guidelines [1] advise the EA designer of a highly complex platform to break the EA down into logical functional areas to enable different parts of the architecture to be implemented with suitable technology. In this way security and safety aspects of the architecture can be 'compartmentalised' and the overall EA designed using cost effective technology for each part of the EA.

Although the functional breakdown can take many forms depending upon technical solution pursued, the following breakdown of logical architectures has been used as a generic case:

- **Command & Control (C2):** Used to connect command and control various sub-systems. This logical bus may be further subdivided as necessary, and includes any backbone databus.
- **Weapon System:** Depending upon the platform role there may be a variety of cannons, missiles and self defence weapons mounted on the platforms. In all cases the weapons system will have safety criticality requirements and need to be certified for safe use. It is likely that at least in part this system will be supplied by the weapons system manufacture and then linked into the main EA electronically through a suitable gateway.
- **Video:** Given the data requirements for transporting digital video around the platform this architecture will require a high bandwidth databus of great flexibility.
- **Survivability:** Depending upon the current mission and platform role there is likely to be at least some aspects of the survivability architecture that will have safety and security implications, especially if a Defensive Aids Suite (DAS) is employed. It is likely that at least in part this system will be supplied by the weapons system manufacture and then linked into the main EA electronically through a suitable gateway.
- **Communications:** A major part of the communications infrastructure will be supplied GFE, an example would be the BOWMAN radio system. This network will be connected to the main EA by a suitable gateway. However there may be other NEC communications channels that require direct integration with the EA.
- **HUMS:** The HUMS architecture will encompass both data generated from LRUs across the EA (e.g. BIT) and data generated by specialist HUMS sensors (e.g. terrain sensor) and a specific HUMS collection point, with a data transfer mechanism to allow various levels of HUMS data to be taken off-platform.
- **Power Management and Distribution:** Although power generation is not directly part of the EA the intelligent management and distribution of the power is, and will require various specialist power switching and data gathering nodes.

In addition to the above in all modern platforms there will be an automotive databus associated with the engine, gearbox and ancillaries. This will be a GFE item as far as the EA designer is concerned, although a gateway will be needed in order to exchange information with the main EA.

Additionally there will be multi-function crewstations where the user can interact with the EA, and where some of the integration of the data from different logical architectures can take place.

The above logical functional breakdown does not necessarily directly relate to a physical technical breakdown as data in more than one of the above logical architectures can be transported on one physical databus, especially where a backbone is employed.

A.2 Examples of the definition of assessment parameters for highly complex EA

A.2.1 Reconfigurability – Adaptability

Reconfiguration relates to changing the platform to meet short-term needs. It is typically characterised by being performed in the field (in theatre), over a short timeframe, using the available skills and tools, and being relatively temporary is easily reversible. Adaptability is focussed on changing the system by direct replacement of system elements.

The guidelines in the VSI metrics list some suggested relevant aspects for assessing the adaptability sub- dimension as:

- Scope: the proportion of relevant changes that can be made in the field (i.e. with available skills and tools, in an acceptable time);
- Difficulty: the work needed to make the changes, in terms of modification to various parts of the system/platform.

It will typically be necessary to define:

- The (kind of) changes that are relevant for a particular assessment;
- The (kind of) modifications that are considered minor.

The scope of the changes to be assessed will be dependent upon components of the logical architecture being assessed, but in general terms this will relate to the physical removal of a LRU and the substitution of a similar or only slightly upgraded substitute. The new LRU should be incorporated into the EA with minimal user intervention.

System Elements – Logical architecture XX	Parameter	Measure
LRU A to be removed	Time taken to remove	X seconds
	Removal Difficulty	Other equipment moved etc
	Tools used	e.g. standard, specialised
LRU B to be installed	Time taken to install	X seconds
	Installation Difficulty	Other equipment moved etc
	Tools used	e.g. standard issue, specialised etc.
Automatic architecture update	Functional check	Check to ensure that full functionality is available with LRU

Table 34: Table for assessing reconfigurability – adaptability sub-dimension

A.2.2 Reconfigurability - Interchangeability

Reconfiguration relates to changing the platform to meet short-term needs. It is typically characterised by being performed in the field, over a short timeframe, using the available skills and tools, and being relatively temporary is easily reversible. Interchangeability is about the number of relevant system elements that can be moved between relevant platform types.

The guidelines in the VSI metrics list some suggested relevant aspects for assessing the interchangeablity as:

- Scope: the proportion of relevant system elements that can be moved;
- Difficulty: the work needed to make the move, in terms of modification to parts of the system/platform.

It will typically be necessary to define:

- The (kind of) system elements that are relevant for a particular assessment;
- The (kind of) modifications that are considered minor;
- The relevant platform types.

An example of the complexity of performing this assessment is the degree of commonality required between different platform variants. The difficulty in moving elements between platforms may differ according to variant, and it is recognised that for some platforms space will be a premium (e.g. a Scout platform). The removal of system elements may be easy for some variants, but there may be difficulty when installing this equipment into another role, due to the lack of space. For vehicle infrastructure elements (e.g. embedded cabling, junction boxes) removal would be extremely difficult even at the platform design authority's premises. Therefore, unless these elements are immediately accessible, it is not expected that these elements would be transferred between platforms.

An example of the desired degree of interchangeability between three main variants of a platform type is given in. It is expected that there for many of the logical architectures there will be some common 'core' LRUs which will be interchangeable, and then some LRUs that are tailored to a specific variant roles. It is expected that the variants in each family will have a very high degree of interchangeability if they are likely to be supplied from a single contractor.

Degree of interchangeability	Varinat 1	Variant 2	Variant 3
Variant 1	High	Medium	Medium
Variant 2	Medium	High	Medium
Variant 3	Medium	Medium	High

Table 35: Table showing the degree of interchangeability between platform variants

For each logical architecture, the assessment will need to assess: the desirable degree of interchangeablity between the variants; the definition of the system element(s) to be interchanged; the time taken to undertake the interchange. A list is required of common LRUs for each sub-variant relating to specific roles.

System Elements – Logical architecture XX	**Parameter**	**Measure**
Main LRU 1	Time taken to remove from platform A	X seconds
	Time taken to install in platform B	X seconds
	Removal Difficulty	Other equipment moved etc
	Installation Difficulty	Other equipment moved etc
	Tools used	e.g. standard, specialised
	Functional check	Check to ensure that full functionality is available with new sensor
LRU 2	Time taken to remove from platform A	X seconds
	Time taken to install in platform B	X seconds
	Removal Difficulty	Other equipment moved etc
	Installation Difficulty	Other equipment moved etc
	Tools used	e.g. standard issue, specialised etc.
	Functional check	Check to ensure that full functionality is available with new LRU

Table 36: Table for assessing reconfigureability – interchangeability sub-dimension

A.2.3 Enhanceability - Capacity

The guidelines in the VSI metrics list some suggested relevant aspects for assessing this dimension as:

- Physical space;
- Power;
- Cooling;
- Bandwidth;
- Operator workload.

For each of these aspects it is suggested that it is necessary to define:

- The capacity aspects;
- The predicted spare capacity needed:
- The basis for the predicted need, including timescale.

For the detailed assessment it is recommended that the aspects list is used apart from Operator Workload, which has many facets not directly within the remit of the EA. It is recommended that the Bandwidth aspect includes latency as the two aspects are both important and are often inter-related.

The predicted spare capacity should be defined to cover a number of different scenarios:

- **Initial capability**: The architecture must comfortably meet all the functional requirements in the SRD, along with the EMC and environmental requirements etc. The relevant parameters are measured and used as a baseline.
- **UOR 5yr timescale**: A significant new capability is added that requires incorporation into the architecture being assessed. This upgrade can be performed in theatre, with assistance from the contractor as required. .
- **Technology Insertion 10yr timescale**: The system will be updated with a step-change in capability. A large part of the architecture being assessed is altered and upgraded. This activity will be undertaken at the contractor's premises.

It is unlikely that a UOR or technology insertion will affect only a single logical architecture, but the assessment will be undertaken on a logical architecture basis. Likewise the assessment of the capacity aspect will be completed for each logical architectural area as in the following table:

XX Architecture – Enhanceability capacity	Current Capacity	Additional capacity for UOR	Capacity after technology insertion
Physical space	Measured	Current + A%	Current + B%
Power	Measured	Current + C%	Current + D%
Cooling	Measured	Current + E%	Current + F%
Bandwidth & Latency	Measured	Current + G%	Current + H%

Table 37; Table for assessing Enhanceability - Capacity dimension sub-dimension

In each logical architectural area the figures for A, B, C, D, E, F, G and H will be defined depending upon the likely future improvements for that logical architecture. These figures could differ between platform variants, depending upon role.

It would be expected that the UOR increase will be significantly lower than that for the technology insertion of most of the logical architectures.

A.2.4 Enhanceability - Modularity

The guidelines in the VSI metrics list some suggested relevant aspects for assessing this dimension as:

- Granularity of modularity expected;
- The way modularity is assessed;

For a high complexity platform assessment the modularity should be assessed on a logical architecture basis. Depending upon the physical implementation of the logical architectures the modularity may be at a physical entity level or software entity level.

Where the level of modularity is a physical entity then this entity may take the form of a complete LRU, or may take the form of sub-system within a physical LRU e.g. a processor card within a rack

For software the level of modularity shall generally be set to reflect discrete functional aspects e.g. distributed interaction with a database or distributed access to a centralised data store.

The assessment shall attempt to relate the expected level of modularity granularity against that produced in the design for the logical architecture. The number of compliant entities, both physical and software, can then be assessed. The table may be extended to include an assessment of each entity.

XX Architecture – Enhanceability modularity	Expected granularity	Number of entities	Number compliant
Level of modularity granularity – physical	e.g. LRU, card, sub-system	X	Y
Level of modularity granularity – software	e.g. functional software entity	A	B

Table 38; Table for assessing Enhanceability – Modularity sub-dimension

At times it may be difficult to quantify the logical architecture with a single level of modularity, e.g. there may one rack mounted element within a wider network of LRUs. In these cases the level of modularity granularity shall be split between the different component parts of the architecture.

A.2.5 Enhanceability - Enablers

The guidelines in the VSI Metrics list some suggested relevant aspects for assessing this dimension as the availability of the skills, tools, knowledge, etc. needed to enhance the architecture.

The parameters are:

- Skills, etc: the extent of skills, knowledge, tools, etc needed to change the architecture;
- Availability: the extent to which the skills etc exist. Skills etc may be available in depth, in a limited way (e.g. just one small team has them), or not at all.

Note: The ease (or otherwise) with which missing skills etc might be obtained is not reflected.

It will typically be necessary to define:

- The relevant skills, etc and those that are considered important;
- The way in which availability is judged, e.g. what constitutes availability in depth or limited availability;
- The scope of availability, e.g. globally, within the UK, within the organisation(s) expected to perform enhancements.

For the detailed architecture assessment, all logical architectures will be assessed individually. This area is a difficult area to judge, being subjective in nature, but it is proposed to subdivide the task into a number of areas to enable a more uniform assessment to be accomplished.

Enhanceability – Enablers Logical Architecture XX	Area to assess	Comment
Skills	Knowledge needed to change architecture	
	Tools needed to change architecture	
	Skills needed to change the architecture	
Availability	Knowledge needed to change architecture	
	Tools needed to change architecture	
	Skills needed to change the architecture	

Table 39; Table for assessing Enhanceabliity – Enablers sub-dimension

A.2.6 Platform Integration - Internal Platform Data Provision

This sub-characteristic relates to the data provision (both sending and receiving data) between the logical architecture under review and the rest of the electronic platform architecture elements.

Some aspects of the electronic interface to other parts of the system will be covered by in great detail by ICDs, (see A.2.15) and the interface to the crewstation (see A.2.8).

The detailed assessment will assess the internal platform data provision on a logical architecture basis. All the data flows and messages from the logical architecture to the other logical architectures, and/or crewstations shall be itemised. Where a gateway or bridge is used to transfer data this should be assessed for its suitability, especially in terms of any buffering system used.

Integration – Internal Data Logical Architecture XX	Data flow	Comment
Transmitted data	Messages to logical architecture 1	List
	Messages to Logical architecture 2	List
Received data	Messages from logical architecture 1	List
	Messages from Logical architecture 2	List
Gateway/Bridge	Messages transmitted	Throughput and Latency should be measured at full load
	Messages received	Throughput and Latency should be measured at full load

Table 40; Table for assessing Integration – Internal Data Provision sub-dimension

A.2.7 Platform Integration - External Platform Data Provision

For all high complexity platforms there will be significant C4I systems requiring data to be both transmitted and received to the wider NEC environment. The boundary between the EA and the supplied C4I equipment may be variable between the different platform roles and the solutions proposed by individual contractors.

It is assumed that for the detailed assessment:

- For a large amount of the NEC data, e.g. including the information to be transmitted by BOWMAN radios, there will be either a single, or very low number of NEC gateways.
- The NEC gateway or communications interfaces are not part of any of the logical architectures.
- Where additional NEC communications devices are accessed directly from a logical architecture, then these should be assessed separately.

The detailed assessment shall check that the relevant information sent from the logical architecture is received at the NEC Gateway and vice versa that the relevant information received by the NEC Gateway is passed to the relevant logical architecture.

Integration – External Data Logical Architecture XX	Messages transmitted by	Messages
Transmitted data to NEC gateway	Messages to NEC	List
Received data from NEC gateway	Messages from NEC	List

Table 41; Table for assessing Integration – Internal Data Provision sub-dimension

A.2.8 Integration – System Control

This dimension addresses how users (human or automated) can control relevant platform resources via the architecture to achieve their objective.

The parameter is

- Scope: the amount of relevant resources that can be controlled in a secure and safe manner under relevant conditions.

It will typically be necessary to define

- the relevant (types of) resources;
- criteria for security and safety, and any other relevant characteristic such as timeliness;
- the conditions of interest;

The resource being controlled may be by a direct link, or may be via another physical architecture. In either case the information being transmitted and received by the logical architecture shall be assessed.

For a high complexity EA platform the number of crewstations per variant (or sub-variant) needs to be identiified by the contractor, an example is given in Table 42.

Platform variant	Platform Sub-variant	User
Variant 1	Sub-variant 1 – section platform	Commander
		Gunner
		Driver
		Crewstation n
		Automatic device 1
		Automatic device 2
Variant 1	Sub-variant 2 – comms	Commander
		Gunner
		Driver
		Comms station
Etc.		

Table 42; Example list of Crewstations for each variant/sub-variant

The assessment of logical architectures will assess the system control required from each User in turn.

Integration – Crewstation Logical Architecture XX	Area to assess	Comment
Data transmitted	Volume	Comprehensive list of messages required.
	Low latency	Specific requirements of low latency messages.
	Security	List of classification of all messages
	Safety related messages	List of all safety related messages
Data received	Volume	Comprehensive list of messages required
	Low latency	Specific requirements of low latency messages
	Security	List of classification of all messages
	Safety related messages	List of all safety related messages

Table 43; Table for assessing Integration– Crewstation Provision sub-dimension

A.2.9 ILS –Built In Test (BIT)

The ILS BIT metric has been included to ensure that the EA can support the transmission of basic BIT data, without specifying the type of functionality of BIT. Many LRUs will support more than one type of BIT, and also transmit other HUMS data.

The guidelines in the VSI metrics list some suggest a parameter for assessing this dimension as:

- Scope: the amount of relevant BIT data supported by the architecture.

It will typically be necessary to define:

- The relevant (types of) BIT data.

For the logical EA assessment the following assumptions are made:

- That all the major electronic system elements, of whatever logical architecture, will support some form of BIT.
- BIT data will be transported and stored across the EA to a single or very limited number of locations.

For every logical architecture the following assessment will be made:

Logical Architecture	BIT implementation assessment
Percentage of major system elements covered	100% of major system elements expected.
Ability to transmit BIT messages	All major system elements should have the ability to transmit BIT messages over at least one physical databus.
Bus Alive Message	Where applicable.

Table 44: Table for assessing ILS - BIT sub-dimension

Note that three types of BIT are possible:

- P-BIT – Power Up BIT performed during vehicle power up;
- C-BIT – Continuous BIT, performed without interrupting normal node operation;
- I-BIT – Interruptive BIT, where a node interrupts its normal operation to perform the BIT test.

The functional requirements of a SRD, along with the technical solutions proposed on individual nodes will dictate the different types of BIT produced, and also the rate for CBIT. This metric only assesses the ability of the EA to support BIT message passing, not the type or functionality of the messages.

Some databus technologies have a built-in bus alive message, often running at a high repetition rate. Where this exists then its functionality should be tested.

A.2.10 ILS – Data Transfer

The guidelines in the VSI metrics suggest a parameter for assessing this dimension as:

- Scope: the amount of relevant data whose transfer is supported by the architecture.

It will typically be necessary to define:

- The relevant (types of) data.

For any platform the data transfer approach will have to comply with the relevant platform HUMS and maintenance requirements. There is an aspiration within the HUMS community to define a common interface for uploading HUMS data, but if no existing interface exists then the design should accommodate the VSI Standards and Guidelines wherever possible.

A number of assumptions are made for the assessment:

- The ILS data will be stored in one or a very limited number of locations;
- The data will be uploaded from the stored location(s) at a rate no more than weekly, but often at a much lower rate;
- Data is no higher than RESTRICTED.

As the relevant HUMS data (to be monitored, stored etc.) is not fully identified the assessment will consist of assessing that the HUMS ICD(s) are fully populated and checking the transfer mechanism and timings:

ILS – Data Transfer HUMS Logical Architecture	Assessment Process
HUMS external ICD	Checked for scope and quality
HUMS interface	Checked to see if VSI recommended standard
Data transfer mechanism	The process – USB stick, direct connection, WI-FI etc. is assessed to ensure feasible. Note that there are HUMS SRD items relating to passing data to adjacent platforms
Time to achieve a week of data transfer	Measured

Table 45; Table for assessing ILS– Data transfer sub-dimension

Note that if parts of an architecture, or logical architectures, are assessed individually then this metric will not always be applicable. When it is not applicable no mark (0) should be assigned.

A.2.11 System Scalability – Vertical Scalability

Vertical scalability is defined as additional resources being available or added to existing system elements to increase their performance. The parameters to assess this will be highly dependent upon the logical architecture and the make up of the nodes within it, but may include:

- Element processing performance: the ability to upgrade the processor or processor board, or add an additional processor;
- Element memory: Ability to upgrade or add RAM/ROM (Random Access Memory/Read Only Memory);
- Storage space: Ability to increase the size of the data stores;
- Software: Ability to update and/or increase the amount of software that can be run. This may require increased memory RAM/ROM.

Even if a system node can be upgraded, the ease of performing this upgrade is a factor - this can be only be addressed subjectively, assessing future upgrade paths cannot be done with any great degree of certainty or accuracy.

Scalability – Vertical Logical Architecture XX	Assessment Process	Method	Comment
Node 1 – general processing node	Increase processing power	Additional processor inserted / Processor upgraded etc.	Assessment of difficulty
	Increase memory	Add additional memory / Add memory or memory module	Assessment of difficulty
	Upgrade of software	Physical process needed to upgrade software	Assessment of difficulty. Any increase in size may result increase memory/processor reqts
Node 2 – data storage facility	Increase in data storage capability	Options available to increase data storage capacity	Difficulty of upgrade

Table 46; Table for assessing System Scalability– Vertical Scalability sub-dimension

A.2.12 System Scalability – Horizontal Scalability

Horizontal scalability is the ability of the system performance to be scaled by the addition or subtraction of system elements. In the case of the logical architectures each must be able to be scaled according to platform role, and allow for future additional elements to be inserted. The changes to the number of system elements should be able to be completed with minimal intervention from the user.

An assessment where a family of platforms is being assessed should take the form of a minimal baseline configuration being defined and then this being scaled to the more complex platform roles, with a future system devised with additional growth nodes added.

Depending upon the logical architecture being assessed there is unlikely to be a requirement for the system to cope with increased numbers of every system element; vertical scalability may be used instead, e.g. a data store is probably more likely to be replaced with one that has additional capacity rather than two of the same size being used.

Scalability – Horizontal Logical Architecture XX	Number of system elements	Comment on change from baseline
Baseline architecture	The number of different elements identified individually	Baseline
Mid-complexity platform variant	The number of different elements identified individually	Comment on procedure(s) for changing over.
Complex platform variant	The number of different elements identified individually	Comment on procedure(s) for changing over.
Future platform variant	The number of different elements identified individually	Comment on procedure(s) for changing over.

Table 47; Table for assessing System Scalability– Vertical Scalability sub-dimension

The tests can be devised to show scalability in the other direction – complex to simple – if required.

Note: There are some elements of overlap between the scalability and capacity sub-dimensions.

A.2.13 Openness – Standards & Technology Selection

Within the VSI Standards and Guidelines lists of recommended standards are given for different aspects of the overall architecture. It is the intention for all new UK MOD platforms that the EA is made up using the recommended VSI standards where possible or open standards where VSI standards are not suitable. A check will be made to ensure that all the technologies used are based upon the VSI recommended or open standards.

It should be noted that the VSI Standards and Guidelines are a 'living' document and so the technology lists will be updated and augmented over time. It would be natural that a future technology insertion in ten years time would take account of newer technologies and that some of the existing standards will have become part of legacy systems.

The assessment will be carried out on a logical architecture basis. The logical architecture will have both the hardware and software aspects checked to ensure recommended open standards are used. There may be a number of entries under hardware and software depending upon the logical architecture being assessed.

Openness – VSI Standards Logical Architecture XX	VSI list Y/N	Comments
Hardware	Y/N	
Software	Y/N	

Table 48; Table for assessing Openness– standards & technology selection sub-dimension

A.2.14 Openness – Documentation

The VSI metrics gives the following advice on the parameters to use:

- Scope: the proportion of the required aspects of the architecture for which high quality documentation exists;
- Accessibility: the extent to which the documentation is available to any organisation likely to need it to enhance the architecture/system.

It will typically be necessary to define:

- Those aspects for which documentation is required (probably 'all' by default);
- What is considered "high quality" documentation;
- The way in which accessibility is assessed. Barriers may include the need to pay a fee or the retention of Intellectual Property. Whether the assessment is purely "as is", or is based on some assumed action such as paying a fee, must be clear.

The scope is each logical architecture. All the information, including data models, should be fully accessible to the MOD, who should also be able to release this information to third parties as required.

Comprehensive documentation should be produced for all aspects of the architecture to enable the benefits of openness to be fully exploited during the life cycle. To have fully open documentation for all aspects of the system may not be desirable for security or safety reasons, but comprehensive documentation should be produced that can be released to approved third parties.

The assessment will be undertaken on a logical architecture basis. There will be a number of documents to be addressed in most logical architecture areas. As an <u>absolute</u> minimum it would be expected that an overview document is produced with supporting documentation to explain hardware and software would be produced. Each document will be addressed on:

- Scope - Is the documentation addressing all the necessary areas?
- Quality – Is the documentation comprehensive and understandable? The areas to be addressed change according to the logical architecture under consideration.

Openness – Documentation Logical Architecture XX	Documentation scope	Documentation quality
Document 1 – Overview		
Document 2 – Hardware		
Document 3 – Software		

Table 49; Table for assessing openness– documentation sub-dimension

A.2.15 Openness – Interface Control Documents

As a consequence of the modularity aspects of the architecture there are defined interfaces both between the logical architecture components and within the logical architecture components. In order for the EA to be augmented or upgraded at a later date fully populated ICDs must be available for release to third parties. These ICDs may be hardware only, software only, or most cases cover both hardware and software i.e. as many of the ISO layers as applicable.

The assessment of this criterion will be on a logical architecture basis. The ICDs will be reviewed for:

- Scope - does the ICD address all the necessary areas?
- Quality – is the ICD fully populated and understandable?

There may be multiple ICDs associated with a particular logical architecture.

Openness – ICDs Logical Architecture XX	**ICD scope**	**ICD quality**
ICD 1 – Hardware		
ICD 2 – Software		
ICD 3 – Hardware/software		

Table 50; Table for assessing openness– ICDs sub-dimension

B Pugh Concept Selection

B.1 Overview

An alternative method of assessing the complexity of the competing tenders rather than combination of scores to give an overall number would be to use the Pugh Concept selection technique. In this technique the strengths and weaknesses of the different tenders can be assessed against a reference datum.

A general view showing the original Pugh concept is shown in Figure 6.

Criteria	Datum (Concept 5)	Concept 1	Concept 2	Concept 3	Concept4
Criteria 1	Datum	S	-	+	S
Criteria 2		S	S	S	+
Criteria 3		-	+	-	+
Criteria 4		+	S	S	-
Criteria 5		-	-	+	-
Criteria 6		-	S	S	S
Criteria 7		S	+	+	+
Number better: Σ+		1	2	3	3
Number worse: Σ–		3	2	1	2
Number same: Σ0		3	3	3	2
Net score		-2	0	2	1
Rank		**4**	**3**	**1**	**2**

Figure 6; Example Pugh Concept Selection diagram

A number of concepts to be evaluated are chosen, in this case five, with one then being assigned as a datum against which the others will be evaluated. The datum ought to be an existing concept, or one that is well understood.

A number of criteria are then chosen, against which the concepts will be evaluated. These criteria will reflect the main needs of the concept, which could have a very wide span, e.g. customer need, ease of manufacture, technical risks.

Each of the concepts is then assessed against the datum and assigned is assigned a 'score' as either: "S" if the concept is similar to the datum; "–" if the concept is worse than the datum; "+" if the concept is better than the datum. In the case of the example diagram the 'score' has been coloured so that the strengths and weaknesses of each concept can be more readily assessed. It should be remembered that there is no weighting of the criteria against each other. This could be added if desired, at the risk of greater complexity.

The number of times each concept is better, worse or the same as the datum can then be counted, an overall score ascribed (number better – number worse), and then the concepts ranked. In this case Concept 1 is very weak and so could be discarded. As Concepts 3 and 4 were better than the datum (Concept 5) then these should be retained for further analysis. Concepts 2 and 5 were effectively the same overall, but if Criteria 7 was thought to be the most important then a subjective choice may be to discard Concept 5 and retain Concept 2. Such subjective decisions, generally taken as part of a team after much discussion, can be part of the process.

It should be noted that at this stage it would be possible to see if combinations of the 'best' of two or more concepts would be feasible.

The process is then repeated picking one of the remaining concepts as a new datum; in this case Concept 3 has the strongest score and is used. The remaining concepts are then scored as before against the new datum. A theoretical output is shown in Figure 7.

Criteria	Datum (Concept 3)	Concept 2	Concept 4
Criteria 1	Datum	-	S
Criteria 2		S	+
Criteria 3		+	+
Criteria 4		-	-
Criteria 5		-	-
Criteria 6		S	S
Criteria 7		-	+
Number better: Σ+		1	3
Number worse: Σ–		4	2
Number same: Σ0		2	2
Net score		-3	1
Rank		**2**	**1**

Figure 7; Example Pugh Concept Selection diagram, 2nd iteration

Examination of the Figure 7 now shows that Concept 2 is worse than the new datum (Concept 3), and so could be discarded. Concept 3 and Concept 4 would seem to be roughly similar and so could be investigated further, either by concept selection iteration or in greater detail using other direct comparison techniques.

Thus the Pugh controlled convergence technique can be thought of in the following steps:

- Definition of the concepts;
- Definition of the selection criteria;
- Compare concepts by means of scoring against one of the concepts which is used as a datum;
- Combinations of the concepts assessed;
- Repetition of the process as necessary
- Assessment of the process to ensure that the result makes sense given the criteria and scoring process.

B.2 Possible Application of Pugh Concept Selection process to multiple EA proposals

To help assess the EAs where a number from different proposals exist, a possible mapping technique would be to use the contractor tenders as the concepts. For a new highly complex platform design the datum for the initial selection iteration would be the either a reference architecture based upon expected scores, or a suitable well understood architecture, perhaps as a result of baselining work. The selection criteria could be either the dimensions or perhaps the logical functional partitions as defined in the VSI guidance documentation. It should be noted that a logical functional decomposition does not infer a physical databus structure; a single physical databus may be used to carry the data for more than one logical functional partition.

An example scoring for four contractors is shown in Figure 9 using a logical architecture breakdown of the EA for a highly complex new platform. The linkage to the VSI metrics remains as the score for each contractors offering in each logical functional area is completed using the standard VSI metrics. The reference architecture would likewise have been assessed using the VSI metrics. The bracketed figure for the reference architecture indicates the error margin for awarding the 'S'(Same) mark i.e 3.8 (0.2) means a score of 3.6 to 4.0 would gain a 'S' mark.

Criteria	Ref Arch	Contractor 1	Contractor 2	Contractor 3	Contractor 4
Command & Control	3.8 (0.2)	3.9	3.6	4.2	3.8
Weapon System	4.1 (0.2)	4.2	3.9	4.2	4.6
Video	4 (0.3)	3.6	4.5	3.5	4.4
HUMS	3.5 (0.3)	4	3.4	3.6	2.8
Communications	3.6 (0.2)	3	3.1	4	3.1
HUMS	3.5 (0.2)	2.8	3.3	3.5	3.5
Power Distribution	3.7 (0.3)	3.4	4.3	4.4	4.6

Figure 8; Example of scoring with four EA contractor's

For assessments of less complex EAs the criteria could be set to the main six dimensions, or all fifteen sub-dimensions if desired.

Once the full assessment matrix has been completed then the 'scores' can be translated to the Pugh selection diagram, as shown in Figure 9.

Criteria	Ref Arch	Contractor 1	Contractor 2	Contractor 3	Contractor 4
Command & Control	Datum Architecture	S	-	+	S
Weapon System		S	S	S	+
Video		-	+	-	+
HUMS		+	S	S	-
Communications		-	-	+	-
HUMS		-	S	S	S
Power Distribution		S	+	+	+
Number better: Σ+		1	2	3	3
Number worse: Σ–		3	2	1	2
Number same: Σ0		3	3	3	2
Net score		-2	0	2	1
Rank		**4**	**3**	**1**	**2**

Figure 9; Example Pugh selection diagram with four EA contractors

Examination of the diagram shows that Contractor 3 and Contractor 4 fare best and have the strongest all-round submissions. These two submissions could then be further assessed by direct comparison or by another iteration of the Pugh concept process in order to find the potentially best solution.

Alternatively the diagram could be used to ensure that only those contractors that have a net score of 0 or above (or any other threshold) are deemed suitable. In this case contractors could be asked to review and re-submit in the areas that they are 'weak' in to further improve their VSI compliance.

Evidence about the ability of any contractor having the ability to meet or exceed the reference architecture can be gleaned from looking across the rows. In the example case one contractor met and three exceeded the reference architecture performance for the power distribution part of the architecture, indicating that this aspect is well understood in industry and mature solutions exist. On the other hand the weapon system architecture had 3 contractors meet the reference architecture performance and one exceed the reference architecture performance, showing that the reference architecture performance is probably somewhere near the limits of current technology. In the case of the communications architecture all were worse than the datum apart from Contractor 3. The reasons for this would need to be investigated further to better understand the reasons why this is so e.g. perhaps one contractor has sole access to some data or information.

A number of issues could arise from this approach, especially arising from the use of reference architecture:

a. The range of scores that generate 'S' need to be defined for each criteria. E.g. if the reference architecture Command & Control score using the VSI metrics was 3.8 then to get a 'S' score a contractor could be 3.8 +/- 0.2. This margin should reflect the error margin in the VSI metrics scoring.

b. It may be that the reference architecture is a 'gold-plated' solution. In this case the contractors would probably mainly score worse than the reference architecture, or occasionally the same. This may result in the first iteration not being informative.

c. Assuming the contractors all have access to the reference architecture it may be that they all meet or mainly exceed the reference architecture requirements – a positive result for assessing overall compliance, but allowing little differentiation between the submissions.

In the case of point a. this may provide the most contentious aspect – the difference between a 'S' and '-' is still rather abrupt and arbitrary. However given a reasonable band for the 'S' score and the fact that the reference architecture will have been accessible to the contractors prior to any submission then this marking scheme is relatively fair.

In the case of point b. the expectations placed in the reference architecture were obviously beyond the current ability of the contractors to comply with. In this event either the reference architecture could be de-scoped by assessing the 'common' aspects of the submissions and using them as the de facto reference architecture, or investigating the reasons behind the submissions e.g. was low cost on over-riding factor

Initial distribution list

External	
Mark Kellaway	FRES IPT
Richard Hall	WS Atkins
Martyn Uttley	DSTL

QinetiQ
Mike Haines (VSI Group)
Bob Connor (VSI Group)
Project File

Report documentation page

Originator's Report Number	QINETIQ/TS/FPPS/TR0900176		
Originator's Name and Location	S J Press, QinetiQ, Farnborough		
Customer Contract Number and Period Covered	Dstl Purchase Order DSTLX-1000007039 Amendment 2		
Customer Sponsor's Post/Name and Location	M Kellaway FRES IPT		
Report Protective Marking and any other markings UNMARKED	Date of issue Aug 2009	Pagination Cover + 59	No. of references 1
Report Title Annex to Vetronics Standards & Guidelines: VSI Metrics for Electronic Architecture Assessment			
Translation / Conference details (if translation give foreign title / if part of conference then give conference particulars) Not Applicable			
Title Protective Marking	UNMARKED		
Authors	S J Press		
Downgrading Statement	Not Applicable		
Secondary Release Limitations	None		
Announcement Limitations	Not applicable		
Keywords / Descriptors	VSI, Metrics		

Abstract

This outlines the VSI metrics to assess electronic architectures designed using the VSI Standards and Guidelines document. Previously the metrics were encapsulated within the VSI Standards and Guidelines document. These have been extracted, updated and put in this stand alone document along with additional guidance.

Abstract Protective Marking:	UNCLASSIFIED

This form meets DRIC-SPEC +SSS issue 7

Blank page