6. ARCHITECTURE EVALUATION

This section presents the evaluation of each candidate framework technical architecture against the architectural evaluation criteria. Subsection 6.1 describes the assumptions made when carrying out the evaluation. Subsection 6.2 reviews the evaluation criteria weighting. Subsections 6.3 through 6.8 present the results of the evaluation for each of the evaluation criteria (cost, implementability, flexibility, manageability, usability, and security). Subsection 6.9 presents a summary of the evaluation results

6.1 Evaluation Assumptions

The following assumptions were made during evaluation of candidate framework technical architectures:

1. Operations and maintenance costs equal three times the system acquisition costs, where acquisition costs equal the total expenses associated with COTS procurement and software development activities. This assumption is based on the premise that 75 percent of total system costs are associated with system operations and maintenance. This premise is supported by Department of Defense and industry studies, which indicate that 60 percent to 80 percent of total system life cycle costs are consumed by operations and maintenance activities[1]. Since the services provided by each candidate architecture are substantially the same, software development costs ate not a significant discriminator among architectures. Therefore, only that portion of the operations and maintenance costs for each candidate architecture associated with the procurement of COTS architectural components (e.g. hardware, system software) is considered as part of the evaluation.

   The calculation in Figure 6-1 shows the derivation of the value for COTS operations and maintenance costs.

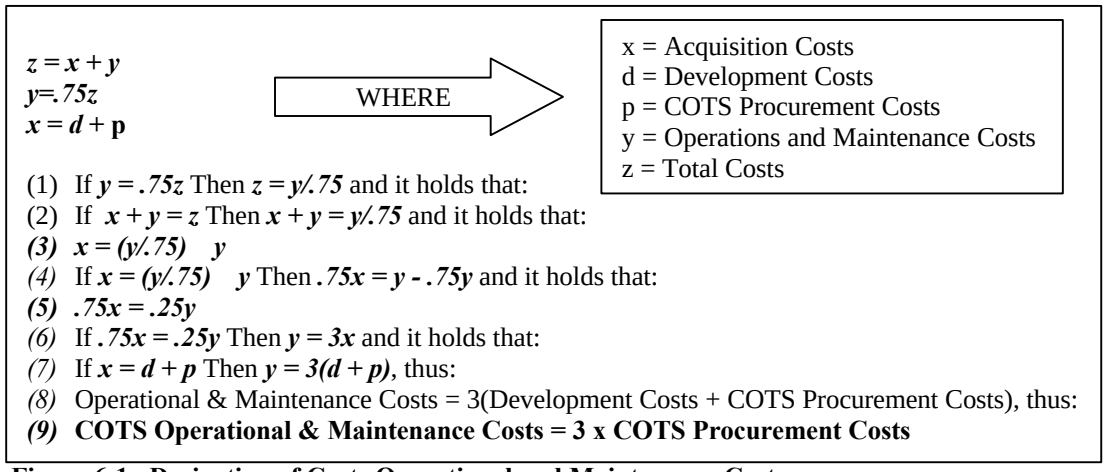

Figure 6-1. Derivation of Costs Operational and Maintenance Costs

[1] For more information refer to the Department of AirForce, Software Technology Support Center's *Guidelines for Successful Acquisition and Management of Software Intensive Systems: Weapons systems, Command and Control Systems, Management Information Systems*

2. Costs associated with system facilities - including heating, ventilation, and air conditioning (HVAC), uninterruptible power supply (UPS), and physical security services - are not significantly different for any of the candidate architectures evaluated. For this reason, these costs are not considered discriminators and are not taken into account during evaluation of the candidate architectures for Project EASI/ED.

3. Costs associated with external system interconnection and networking services, including subscriptions to Value Added Network (VANs) and Internet Service Provides (ISP) are not significantly different for any of the candidate architectures being evaluated. That is, all of the candidate architectures will be similarly networked with external trading partners and users. For this reason, these costs are not considered discriminators and are not taken into account during evaluation of the candidate architectures for Project EASI/ED.

4. All architectures are implemented at a geographically centralized location. As a result, costs associated with internal system interconnection and networking services, such as those associated with distributed system component communication, will be not be significantly different for any of the candidate architectures being evaluated. For this reason, these costs are not considered discriminators and not taken into account during evaluation of the candidate architectures for Project EASI/ED.

6.2 Evaluation Criteria Weighting

Weights for the evaluation criteria were determined using the AHP methodology in consultation with ED staff. This process is described in detail in subsection 4.2. Figure 6-2 shows the weight assigned to each evaluation criterion.

CRITERION	Relative Importance
Usability	20%
Cost	15%
Manageability	20%
Security	15%
Implementability	20%
Flexibility	10%

Figure 6-2. Weights of Evaluation Criteria

6.3 Cost Evaluation

The cost criterion measures expense associated with architecture capital investments and operations/maintenance activities. This criterion comprises 15 percent of the final evaluation score.

The cost evaluation is based on the estimated costs summarized in Figure 6-4 . These costs were calculated based on the prices of hardware and software components that populate each candidate architecture. Refer to Appendix H for (a) a full list of vendor products used in the architectures and the costs of each, and (b) for a build up of the estimated costs associated with each candidate framework architecture.

The same ratio of operations and maintenance costs to acquisition costs is used for all four candidate architectures. In general, experience has shown that the costs associated with operating distributed (e.g. client/server based) systems are higher than the costs associated with operating centralized systems (e.g. mainframe based). However, a large majority of these extra costs are due to increased requirements for end-user PC support. Schools and students make up the majority of the user community for Project EASI/ED, and PC support for these users is not expected to be a significantly cost factor for ED.

The distributed processing/centralized data candidate architecture is substantially less expensive than the other candidate architectures because it lacks the dedicated decision support, data warehousing, and database replication hardware and software that the "replicated data" candidate architectures include. The centralized processing/centralized data candidate architecture does not have these components either, but a fully centralized solution requires a much larger server (the Ultra Enterprise 10000) than the other candidate architectures. This somewhat offsets the cost differential. The distributed processing/replicated data for publication candidate architecture is a little cheaper than the distributed processing/replicated data for consolidation candidate architecture because it requires only one transaction processing database server – the distributed processing/replicated data for consolidation candidate architecture requires two.

Figure 6-3 shows the relative strength rating given to each candidate architecture for the implementability criterion. This strength rating was calculated by awarding a score to each candidate architecture in comparison with the other candidate architectures, and then applying the AHP methodology described in subsection 4.2. Appendix D lists the scores awarded to each candidate architecture.

	Relative Strength Rating
Centralized Processing/Centralized Data	22%
Distributed Processing/Centralized Data	53%
Distributed Processing/Replicated Data Publication	15%
Distributed Processing/Replicated Data Consolidation	11%

Figure 6-3. Strength of Candidate Project EASI/ED Architectures for Cost Criterion

Cost Category	Centralized Processing / Centralized Data	Distributed Processing/ Centralized Processing	Distributed Processing/ Replicatied Data for Publication	Distributed Processing/ Replicatied Data for Consolidation
Software	$2,151,232	$2,848,124	$4,284,472	$4,403,760
Hardware	$10,494,440	$6,132,795	$9,067,940	$9,103,045
Total Acquisition Cost	$12,645,672	$8,980,919	$13,352,412	$13,506,805
Operations and Maintenance	$37,937,016	$26,942,757	$40,057,236	$40,520,415
Total Cost	$50,582,688	$35,923,676	$53,409,648	$54,027,220

Figure 6-4. Project EASI/ED Candidate Architecture Costs

6.4 Implementability Evaluation

The implementability criterion measures the degree to which technologies comprising the architecture are mature, understandable, COTS-based, and supportable by available skilled personnel. This criterion comprises 20 percent of the final evaluation score.

The fully centralized candidate architecture evaluated well on this criterion, primarily because this is the most mature and easily understandable architecture amongst the candidates. Availability of skilled personnel is also relatively high for the fully centralized architecture, both because of the maturity of the technology and because skills in distributed system management tools are not required. The use of similar hardware and software components for all candidate architectures mitigated the factor of the use of COTS solutions. However, operating systems traditionally viewed as supporting a more centralized architecture (e.g. IBM/MVS) do lag behind traditionally decentralized operating systems (e.g. UNIX) in availability of the newer COTS applications based on portable databases (e.g. Oracle, Informix, or Sybase).

Each of the decentralized candidate architectures received the same evaluation score for implementability because they basically share the same strengths and weaknesses. Tools for decentralized system management are less mature than those for centralized systems. The availability of personnel skilled in the use of these tools is correspondingly less. There is a wide range of COTS-based applications available for these architectures, especially the newer applications that are based on client-server and Internet technologies.

Figure 6-5 shows the relative strength rating given to each candidate architecture for the implementability criterion. This strength rating was calculated by awarding a score to each candidate architecture in comparison with the other candidate architectures, and then applying the AHP methodology described in subsection 4.2. Appendix D lists the scores awarded to each candidate architecture.

	Relative Strength Rating
Centralized Processing/Centralized Data	40%
Distributed Processing/Centralized Data	20%
Distributed Processing/Replicated Data Publication	20%
Distributed Processing/Replicated Data Consolidation	20%

Figure 6-5. Strength Rating of Candidate Project EASI/ED Architectures for Implementability Criterion

6.5 Flexibility Evaluation

The flexibility criterion measures the degree to which architecture components are open to product/vendor heterogeneity, based on widely accepted standards, and scalable. This criterion comprises 10 percent of the final evaluation score.

The fully centralized candidate architecture scored low in the “openness” measure because a centralized solution (where the majority of functionality is co-located on a single hardware component) offers much less flexibility in replacing system components. Interchanging any one part of the system can dramatically impact all other system functions. Conversely, if an architecture is distributed, changes to one system component can be more easily shielded from the rest of the system.

Distributed architecture components are also more widely interchangeable across because they are more likely to communicate with each other using industry-standard interfaces. If components use standards widely accepted by information technology vendors, then there is the potential for the replacement of a system component from one vendor with an equivalent one from another vendor to be largely transparent to the rest of the system. System components that are operating on a single, centralized hardware platform are more likely to require replacement with a product from the same vendor that supplied the platform.

The most significant difference between the "distributed processing/replicated data" candidate architectures and the "centralized data" candidate architectures is in the area of scalability. An architecture is considered scalable if it can accommodate a wide variety of system sizes, in terms of both number of users and volume of data, and still provide acceptable performance. "Distributed processing/replicated data" architectures are inherently more scalable because they offer the capability to add more application/data server components as demands on the system rise. Centralized architectures are constrained by the limit on how many processors/disks can be accommodated by the central processing resource. The "distributed processing/centralized data" architecture evaluates better that the fully centralized architecture because it has the flexibility to increase the number of application servers it contains. It is, however, still limited by a single central data resource. Since both transaction processing and decision support processing must occur on the same database, the performance of one or the other (or both) will likely be compromised.

Figure 6-6 shows the relative strength rating given to each candidate architecture for the flexibility criterion. This strength rating was calculated by awarding a score to each candidate architecture in comparison with the other candidate architectures, and then applying the AHP methodology described in subsection 4.2. Appendix D lists the scores awarded to each candidate architecture.

	Relative Strength Rating
Centralized Processing/Centralized Data	11%
Distributed Processing/Centralized Data	16%
Distributed Processing/Replicated Data Publication	37%
Distributed Processing/Replicated Data Consolidation	37%

Figure 6-6. Strength Rating of Candidate Project EASI/ED Architectures for Flexibility Criterion

6.6 Manageability Evaluation

The manageability criterion measures the degree to which the technologies comprising the architecture are reliable, available, serviceable, and controllable. This criterion comprises 20 percent of the final evaluation score.

All four candidate architectures evaluated equally well on the manageability criterion. This was for the following reasons:

- Centralized architectures are generally easier to control than distributed architectures, because all system resources are operating on a single, central hardware platform. There are fewer of the complexities of dealing with products from multiple vendors and the telecommunications necessary to link these products. Centralized architectures are traditionally reputed to be reliable and available, although they suffer an element of risk due to their dependence on one central resource. This central resource, if it does fail, can cause a catastrophic failure for the system.

- Distributed architectures spread the risk of system failure across many system components, so that a single component's failure can be isolated and have minimal effect on the rest of the system. Also, distributed architectures offer the possibility of clustering system components so that each component benefits from full redundancy and failures become almost transparent to the user. Distributed architectures, however, are generally viewed with reason as more complex to control than centralized architectures. The centralized software management and dynamic deployment strategy (see subsection 5.3.4) used by each of the distributed candidate architectures goes some way to alleviating those complexities.

Figure 6-7 shows the relative strength rating given to each candidate architecture for the manageability criterion. This strength rating was calculated by awarding a score to each candidate architecture in comparison with the other candidate architectures, and then applying the AHP methodology described in subsection 4.2. Appendix B lists the scores awarded to each candidate architecture.

	Relative Strength Rating
Centralized Processing/Centralized Data	25%
Distributed Processing/Centralized Data	25%
Distributed Processing/Replicated Data Publication	25%
Distributed Processing/Replicated Data Consolidation	25%

Figure 6-7. Strength Rating of Candidate Project EASI/ED Architectures for Manageability Criterion

6.7 Usability Evaluation

The usability criterion measures the degree to which the architecture improves system and data usability, while masking system complexities. This criterion comprises 20 percent of the final evaluation total.

The centralized process/centralized data and distributed process/centralized data candidate architectures evaluated poorly for the usability criterion. This is because of limitations that the centralized data architecture places on users. In a centralized data architecture, all database operations are acting against the same, centrally located, data resource. This means that both transaction processing and decision support applications are using the same database. Not only could this cause performance problems, with contention for database resources slowing down response times for users, but specialized data warehousing facilities that require multi-dimensional database technology are not available.

The distributed processing/replicated data for publication and distributed processing/replicated data for consolidation candidate architectures both offer data warehousing functionality, and separate transaction processing and decision support processing across different databases. The “replication for consolidation” candidate architecture is evaluated less strongly than the “replication for publication” candidate architecture because the increased complexity involved in synchronizing multiple databases being updated by multiple users is difficult to disguise from users.

Figure 6-8 shows the relative strength rating given to each candidate architecture for the usability criterion. This strength rating was calculated by awarding a score to each candidate architecture in comparison with the other candidate architectures, and then applying the AHP methodology described in subsection 4.2. Appendix B lists the scores awarded to each candidate architecture.

	Relative Strength Rating
Centralized Processing/Centralized Data	12%
Distributed Processing/Centralized Data	12%
Distributed Processing/Replicated Data Publication	44%
Distributed Processing/Replicated Data Consolidation	31%

Figure 6-8. Strength Rating of Candidate Project EASI/ED Architectures for Usability Criterion

6.8 Security Criterion Evaluation

The security criterion measures the degree to which the architecture provides adequate authentication, information confidentiality and integrity, access control, security administration, and auditing services, as justified by business needs. This criterion comprises 15 percent of the final evaluation score.

All the candidate architectures described in the *TVTA Report* provide a high measure of application security, implementing multilevel security solutions that provide authentication, access control and confidentiality services. However, there is one area in which the fully centralized candidate architecture is much less secure than the three distributed candidate architectures. This is in the area of Internet access security. The three distributed candidate architectures place a Web server outside a "firewall", where it is exposed to the Internet but separated from the other system components by the security features provided by the firewall. If the security of the Web server is "breached", the web server can be recycled and dynamically reconfigured by the system – quickly correcting any unwarranted system modification. Firewalls are also capable of maintaining detail audits logs of all transactions crossing the firewall. Within the fully centralized candidate architecture all system components are physically co-located on the same central hardware platform – including the web server software. This makes it impossible to place a firewall between the Web server software and the rest of the system.

Figure 6-9 shows the relative strength rating given to each candidate framework technical architecture for the usability criterion. This strength rating was calculated by awarding a score to each candidate architecture in comparison with the other candidate architectures, and then applying the AHP methodology described in subsection 4.2. Appendix B lists the scores awarded to each candidate architecture.

	Relative Strength Rating
Centralized Processing/Centralized Data	8%
Distributed Processing/Centralized Data	31%
Distributed Processing/Replicated Data Publication	31%
Distributed Processing/Replicated Data Consolidation	31%

Figure 6-9. Strength Rating of Candidate Project EASI/ED Architectures for Security Criterion

6.9 Architecture Evaluation Summary

This section summarizes the Project EASI/ED candidate architecture evaluation results. Figure 6-10 shows abbreviations for each of the candidate architectures that are used as the column headings in the following figures. Figure 6-11 summarizes the relative strength ratings given to each candidate architecture for each of the evaluation criteria. The percentage figures used in Figures 6-3 to 6-9 were converted to decimals and summed to give an unweighted total for each candidate architecture. This is shown in Figure 6-11. Using the AHP methodology described in subsection 4.2, the evaluation criteria weights shown in Figure 6-2 were applied to each of the scores for each of the candidate architectures. The result of applying these weights is shown in Figure 6-12, which presents the final evaluation results.

Architecture	Abbreviation
Centralized Processing/Centralized Data	CP/CD
Distributed Processing/Centralized Data	DP/CD
Distributed Processing/Replicated Data Publication	DP/RDP
Distributed Processing/Replicated Data Consolidation	DP/RDC

Figure 6-10. Candidate Project EASI/ED Architecture Acronyms

Criteria	CP/CD	DP/CD	DP/RDP	DP/RDC
Cost	0.22	0.53	0.15	0.11
Implementability	0.40	0.20	0.20	0.20
Flexibility	0.11	0.16	0.37	0.37
Manageability	0.25	0.25	0.25	0.25
Usability	0.12	0.12	0.44	0.31
Security	0.08	0.31	0.31	0.31
Total	**1.18**	**1.56**	**1.71**	**1.54**

Figure 6-11. Unweighted Candidate Project EASI/ED Architecture Evaluation Scores

Criteria	CP/CD	DP/CD	DP/RDP	DP/RDC
Cost	0.03	0.08	0.02	0.02
Implementability	0.08	0.04	0.04	0.04
Flexibility	0.01	0.02	0.04	0.04
Manageability	0.05	0.05	0.05	0.05
Usability	0.02	0.02	0.09	0.06
Security	0.01	0.05	0.05	0.05
Total	**0.21**	**0.26**	**0.28**	**0.25**

Figure 6-12. Weighted Candidate Project EASI/ED Architecture Evaluation Scores