निर्वाचन आयोग, नेपाल
ELECTION COMMISSION, NEPAL

The Voter List Applications

Findings and recommendations on
the software architecture and development

Morten Forfang
December 2010
Version 1.0

UNDP Nepal

Table of Contents

Please consult the dictionary, section 5.1, for particular abbreviations and terms.

1. Executive summary

This is a summary of the preliminary findings (cf. section 4):

- An **application layer** needs to be introduced to the CVLA (Central Voter List Application). This is a software layer between the database and the clients. It serves to provide data independence as well as a more sustainable way to provide services, security, accountability and maintainability. An application server product typically supports the application layer.
- It is recommended that the CVLA has an appropriate design and functionality for dealing with different **external software systems** including the DVLA (District Voting List Application) and the GIDC/NID (Gov. Integrated Data Center/National ID). The functions will be implemented in the application layer, and will serve to insulate the internal structure of the VR (Voting Registry) and provide targeted services for the various consumers of VR data. The functionality is realized as custom developed code employing services from the application server.
- It is recommended that the CVLA has an enhanced design and functionality to deal with user operator **identity, access, accountability and audits**. The present 3 VR applications have reasonable security and accountability mechanisms built in, but are probably too primitive to deal with a general, maintainable solution catering for a more sustained identity model and a richer set of permissions and roles. These functions can be realized as separate products or as custom developed code running in the main application server or independently.
- The agency that has the **operational responsibility** of the CVLA needs to be a competent, capable body that can handle application operations in a standardized, best-practices way. My impression is that the EC IT doesn't have the necessary capability or competence to fulfill that role as it stands today. Steps are needed to make sure that either the capacity of the EC IT is strengthened or that some other organization assumes this role.
- The organizational setup for the **maintenance** of the CVLA, needs to be such that both the developer and the CVLA owner have experience and competence with running software development projects.
- My impression is that the **EC IT** doesn't have the necessary capability or competence to fulfill the client role as it stands today. Steps are needed to make sure that either the capacity of the EC IT is strengthened or that some other organization assumes this role.
- It is unclear whether the **Vendor** has the necessary capability or competence to fulfill the developer role as it stands today. Further investigation is warranted to make a reasonable recommendation for this role.
- It is recommended that the CVLA is under a sustained, well anchored, competent **IT governance** body with decision-making authority. My impression is that the EC IT doesn't have the necessary capability or competence to fulfill that role as it stands today.

Steps are needed to make sure that either the capacity of the EC IT is strengthened or that some other organization assumes this role.

- It is recommended that steps are taken to ensure **sufficient quality and trustworthiness** of the voter registration data presently collected. Data accountability is presently not strong enough. It is unclear whether the voter entry data are of sufficient quality. Management of software and database releases, incoming data from the districts, and rollout of new software need to be managed in a systematic and transparent fashion.

2. Introduction

I have over the period December 3rd - 12th, 2010, interviewed relevant parties at the UNDP ESP, EC IT, Worldlink, GIDC, NID and the HLCIT. The work was done at the request of UNDP ESP.

2.1 Aim and scope of work

The aim has been to make recommendations for the improvement of the software development endeavor in general and specifically to improve the software architecture for the CVLA. My main focus has been to make recommendations so that there is an increased chance the software solution is maintainable and that the data remain trustworthy.

The recommendations, found in section 9on page 4, are generally based on a set of findings, found in section 3 below.

2.2 Acknowledgements

The staff at the EC IT and the UNDP ESP has been friendly, supportive and in general it would have been impossible to make findings and recommendations without them. In particular I acknowledge the people referred to in section 5.3, who have been actively contributing to the subject-matter and helped, shape this report.

3. Findings

3.1 UNDP

There are UNDP Documents that specifically relate to the three VR applications;

- W.r.t. the requirements, specification, design and development of the three VR applications,, I have found three documents that seem particularly relevant [1][2][3]. They are all authored by UNDP Nepal and cover the VR application software requirements, the choice of databases and the integration landscape over time.
- W.r.t. plans, milestones, and acitivities there are two documents that seem relevant [4][5]. Both are authored by UNDP Nepal. These plans are very general and of coarse granularity.

- W.r.t. budgetary matters, there seems to be two documents that seem relevant [6][3 Annex H]. One covers the entire Voter List with Photo project, whereas the other is a detailed look at the database costs, including hardware and the physical building.

3.2 EC IT

From continuous conversations and interviews Dec. 2-5th 2010. Twice reviewed with EC IT staff.

- The 3 VR applications - techincal considerations;
  - The DCA is a .net single-user 2-layer application consisting of a rich client connected to a local oracle database.
    - It features a login with password authentication. In practice a common login is frequently used across many laptops.
    - The database is encrypted with a common key. The data are not signed and there is no link between operator/logged in user and encryption. Each record has a field for which operator did the registration.
    - The database comprises all multimedia contents.
    - Export is possible that includes the three pictures for each record. An option to export with the face pictures copied into a file folder exists and is there as a consequence of a NID (presumably GIDC) requirement.
    - The face picture acquisition features an automatic normalization function, is usable and working well.
    - The signature acquisition seems awkward and hard do in a consistent manner.
    - The SRS requires functionality for database import, reporting and having secure audit trails. These I have not been able to verify exists. *Update 101214:* I have by email received a sample audit log from Worldlink.
    - There is work going on to update the application w.r.t. allowing multiple forms of voter authentication.
    - A crucial feature is that each voter receives a unique id – a registration number. The registration number stems from a physical form with a printed/stamped number. My understanding is that this serves as a key in the DCA/DVLA/CVLA databases. The uniqueness contraint is dependent on there over years being a systematic numbering system in place at the ECN.
  - The fingerprint acquisition features at DCA depend on an appropriate thresholding to be useful for matching at CVLA. The roundtrip from gathering data from DCA, testing the fingerprint templates at CVLA and then possibly recalibrating the DCA fingerprint acquisition thresholds is slow and cumbersome.
  - Pending an actual look at the information model employed, I find no indications that records are versioned (or a history is stored). Thus, for a continuous update feature, the previous version of a record will be lost.

- There is at the EC no installation of the DVLA or CVLA. The development status w.r.t. to these at the vendor is unclear, although I'm told that the DVLA has been received by the EC IT but not tried or used. After several rounds with the EC IT it appears staff cannot find the install files for the DVLA. Only one staff possesses a copy it appears. *Update 101215:* Wordlink was at EC IT installing the DVLA and delivering source and install files. The plan was to install the CVLA as well, although licensing matters with the Megamatcher AFIS were expected to make it non-functional. Whether what was delivered is sufficient to do a full separate installation by non-worldlink personell remains to be seen. *Update 101216:* The CVLA was not installed. Nor were the necessary database scripts delivered for the DVLA/CVLA to be able to do a separate installation.
- The scope of the vendor's work is all three VR applications. Maintenance responsibility is also the vendor's. It is unclear over what timespan this maintenance responsibility runs. There is no formal agreement or contract stating the vendor has responsibility for maintenance of the applications, but there is an understanding between the parties, if I understand it right. The contractual aspects of the relation w.r.t. development remains unclear. All contractual documents I'm told are in Nepali, so I'm not able to verify the contractual relation governing the relation between the EC IT and the Vendor. *Update 101215:* Sachin were able to find a copy of the contract [12] with the Vendor. It was all in English. The scope is indeed all three VR applications, including continuous registration. There is a 1 year support provision in the contract as well.
- The EC IT has a user guide and an installation guide for the DCA. Apart from that, there exists no documentation regarding the vendors work. No technical documentation, no architecture, no design, no information model. Such documents will not be forthcoming to the EC IT before the vendor's work is completed. *Update 101215:* There are a number of specified required documents to be produced in the contract [12]. Wordllink has sent us some documents, including E-R diagrams.
- The EC IT has no formal/printable plan for their work on either of the VR applications. There are no available/printable milestones, activities, schedules or plans governing the work (see the UNDP findings for more general, project-wide plans). It is unclear when final delivery of any of the applications are planned for.
- The EC IT plans to have the operational responsibility for all VR applications themselves, including the CVLA. They are 7 people, including 3 temporary staff and 1 from undp.
- The EC IT development methodology remains a bit unclear but my impression is that it is possibly based on a phased approach, but in practice works like a waterfall method. Very non-agile and offhand w.r.t. the vendor.

3.3 The Vendor - Worldlink

From a roughly two-hour meeting and demo with Rojina (rojina@wlinktech.com) and an AFIS expert at the Vendor's premises on Dec. 5th, 2010.

- Rojina states all three applications have been delivered and demoed to/for the EC IT.
- The CVLA is a hybrid (i) two-layer rich client on top of a database and a (ii) 3-layer browser-web server-database in combination. The latter part is for reporting purposes, the former for all other types of functionality.
- (i) is essentially an expanded version of the DVLA.
- A CVLA demo was performed of (i).

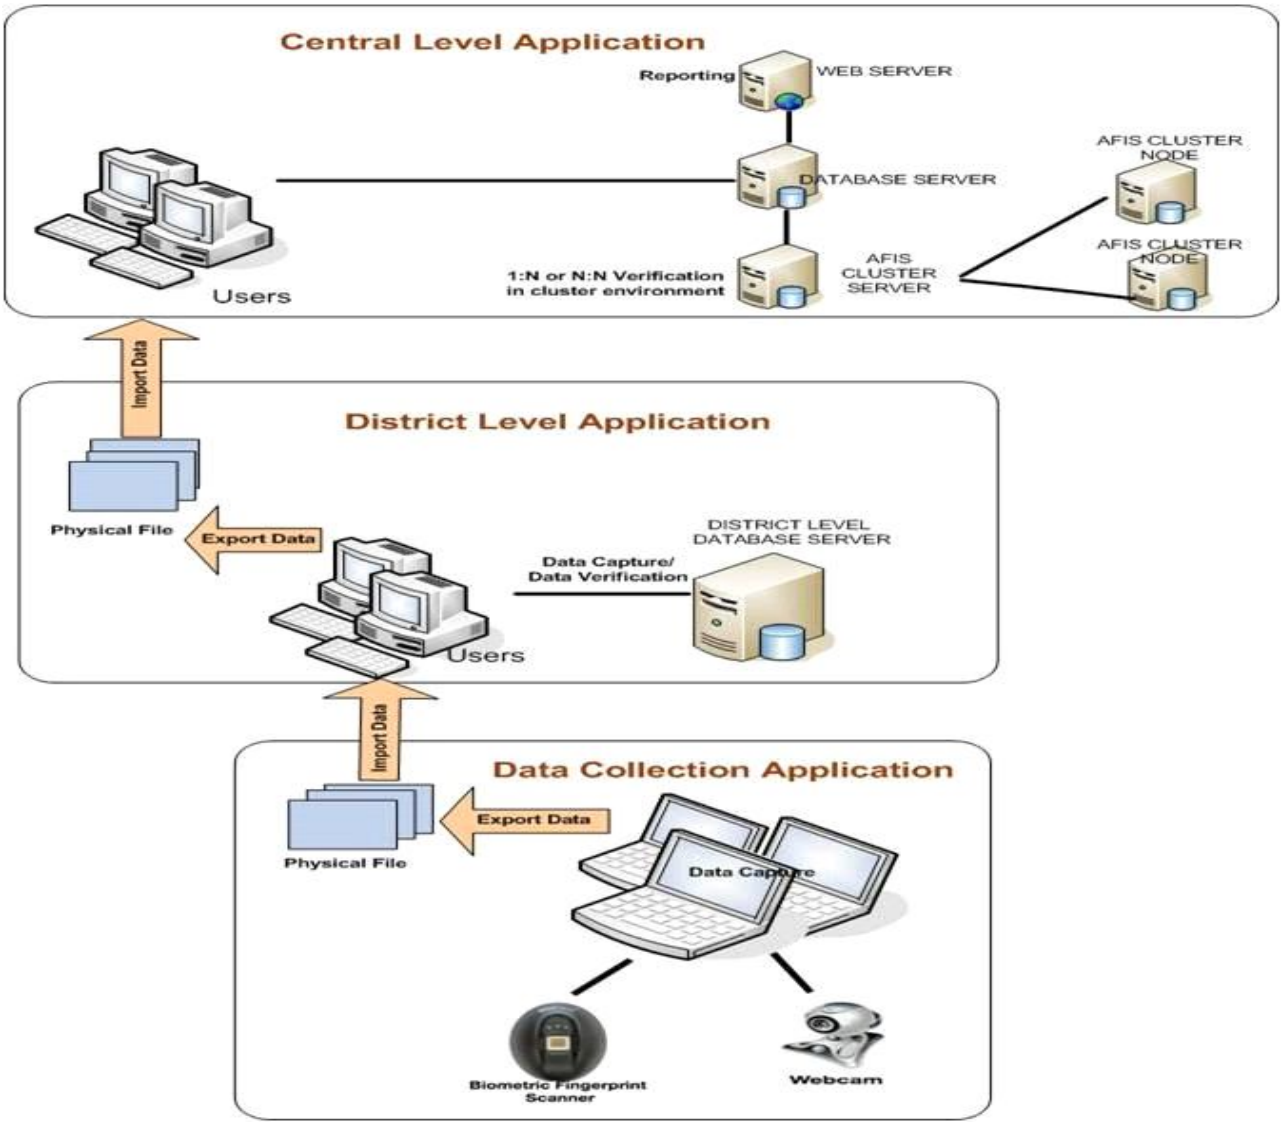

Figure 1: The vendor's CVLA architecture

- There seems to be a minimum of documentation underpinning the work, including a rudimentary application architectural sketch and a behavioural diagram (overall

workflow). There is apparently an information model (an E-R diagram), but it was not available. I've requested that by email. *Update 101214:* the documentation from the meeting has been supplied.

- If I understand it rightly, the business rule for how to merge two records with the same registration number (e.g. an edited version from the DVLA is imported into the CVLA which already has a record with the same reg. number) is to keep the oldest and discard the newest. Apparently this rule stems from a specific request to the EC (IT?). *Update 101215:* If I understood Ram rightly today, there are different business rules for DCA->DVLA imports and DVLA->CVLA imports. For the former, only new records are to be accepted (continuous registration is not to be undertaken by the DCA) and thus, any new record with the same reg. number (form number) as an old one, will be discarded. For the latter, the rule is that if two records have the same reg.number, the oldest will be discarded.
- The security of the database contents is reasonable. Circumventing the protection seems quite possible, but would require a reasonable amount of technical insight.
- The accountability of the data is also reasonable. The design is such that one in most cases would be able to trace a record change back to a specific user. It is unclear however over exactly what scope (what type of operations) the accountability features work on. There are likely to be loopholes when one takes into account the importing and exporting of data and users back and from DCA-DVLA-CVLA, such that data modifications have been made without one able to trace it back to somebody identifiable and accountable. *Update 101215:* Audit traces are discarded both for DCA->DVLA imports and DVLA->CVLA imports.
- The vendor seemed frustrated with the lack of detailed specifications and the inability to make a delivery, get it accepted and move on.
- It is clearly hard to gauge the quality and methodology of a software development effort from a meeting and a demo. In general the vendor seems capable of making two-layer persistent applications with acceptable technical quality and seems quite capable of making tactical technical decisions in absence of detailed specifications.
- The Vendor's capacity to reasonably manage releases, changes and configurations suitable for sustained development and maintenance of the CVLA remains unclear.

3.4 The GIDC

From a half day visit, including a meeting and a tour of the premises. We met Deputy Directory Bikal Paudel (bikal@nitc.gov.np), Assistant Director Sudeep Dangi (sudeep@nitc,gov.np) and Computer Engineer Sunil Paudel (sunilpaudel@nitc.gov.np) among others.

- The Gov't Information Data Center is under the National Information Technology Center (NITC). The NITC is under the Ministry of Science and Technology.

- It is presently a building fir for hosting servers, a large server with control room. There are cooling (?), fire and cabling facilities. There was ample physical free capacity it seems.
- The personell seemed reasonably competent on physical hardware and network. We were shown network monitoring software. My impression is that staff had little capacity to be running application operations.
- They work with the HLCIT on Govt. EA.
- It seems most prospective partners (ministries) think of the GIDC as a place to have backups. This I understand is because the connectivity outside the largish campus they have is limited. Ministries do in general not want their primary servers at the GIDC.

3.5 The NID

From a two hour meeting at their premises with deputy Prahalad Pokhrel and Rajendra Sigdel.

- Generally on the same page as the progression plan [5]. First they will use the CVLA, and then they will use their own National ID Data Centre.
- The NIDC will get its data in a one-off transfer. Didn't have the presence of mind to ask when they would start printing, but I guess it had to be after VR reg completion since it's a one-off.
- They envisage an update process using their district offices. Canham tells me they don't have a ward/local level though, so they would struggle. The later stages of [5] plans to use VDCs which are offices at the local level and which sort under Min. of Local Development. The PC's bought match roughly the number of VDC's such that they can be taken over by them.
- They had quite a number of good looking diagrams on application architecture, network topology etc.
- They are planning to have their own staff and run their own registry, but use GIDC as a secondary site and an off-site backup site.
- An ADB funded technical specialist is apparently coming in January to do planning, paving the way for procurement.
- NID card fields;
    - ID number (15 digits, 8 from birthdate and 2 from region - a structured number, cf. Steve's document), Name in nepali and english, place of birth, perm. address, photo + ghost photo (uncopyable), signature, expiry date (~10-15 yrs), chip?, issued at and by, MRZ (?) machine readable field, fingerprint(s?). No ethnicity, religion. Civil status may be stored in chip.
    - So there are at least 3 pictures on this card.. a lot?

3.6 The HLCIT

From a two hour visit to the High Level Commission for Information Technology on Dec. 12th, 2010 with Thapa and Sudip. We met with Manohar Bhattarai (manohar_kb@hlcit.gov.np), the Vice Chairman, Juddha Gurung (briefly), a member secretary, a project coordinator I'm lacking the name of and two consultants from PriceWaterhouseCoopers India; Joydeep Chakraborty (joydeep.chakraborty@in.pwc.com) and Goutam Rath (goutam.rath@in.pwc.com).

- The PWC consultants sketched a SOA/Mule based central EA for the govt. of Nepal. Mentioned counterparts were the EC, dept. of transport, a public service commission, a series of service centers, Tax/Revenue dept and a national portal. The central ESB is planned to be hosted by the GIDC.
- These service centers are out in the field and may one day be capable of offering electronic services to the public
- The PWC consultants have come quite far in producing an Govt. EA report/guide (?) which has sections on Data, Security, Infrastructure and Applications. On request, they volunteered to send me a copy of the report although it wasn't finished. *Update 101215:* We've received 12 Enterprise documents for the Govt. of Nepal [13] from the PWC consultants. Very large and comprehensive.
- The HLCIT seemed quite familiar with the thought that the EC VR may become the de facto source of citizen information for the foreseeable future. They were also familiar with the thought that ideally the VR would be transferred to the NID for them to manage.
- The PWC consultants seemed quite technical and focused on deploying a pilot onto their ESB solution at the GIDC. Capacity building, training and IT Governance seemed to come as an afterthough and because I asked. Apparently there is being written a governance document as well, but if I understood it rightly, I couldn't see that work before it had been approved or seen to by the HLCIT.
- Central to the HLCIT is e-government and Manohar admitted it seemed to be more about organizational change and BPR and less about technology. Apparently Manohar is a computer engineer by background. The e-government effort is funded by the ADB, if I understood it correctly. The PWC presence is for a five year period.
- I asked about there being planned a service registry and Joydeep stated it would be taken care of by LDAP (?).
- Identity is to be managed by an LDAP service. It remains unclear as to how organizational boundaries are to be treated, as I didn't really understand the PWC answers to my questions on nation wide identity management (in case they're planning on that) or federation issues.
- Indian PWC may be an interesting agency to run the CVLA operations.
- *In conclusion the target EA may make sense, is very ambitious and all-encompassing and my impression is that there doesn't seem to be too strong an effort on capacity building.*

4 Recommendations

4.1 Recommendation 1 - Application Layer

4.1.1 Summary

Recommendation: The CVLA is recommended to be set up as, at a minimum, three layer server-client application.

Rationale: In addition to a database layer and a set of clients, there should be an application layer that serves to cater for business logic, lessen coupling with external systems, manage identity, authorization and access as well as providing enhanced auditing and logging capabilities. This is a basic, very common architectural blueprint for enhancing system maintainability, security and accountability.

Technical details: The architecture should to the extent possible rely on standards and industry-wide blueprints, like those governed by w3.org and www.oasis-open.org, www.enterpriseintegrationpatterns.com as well as platform blueprints for Miscrosoft technology (http://msdn.microsoft.com/en-us/practices, http://msdn.microsoft.com/en-us/architecture, http://msdn.microsoft.com/en-us/library/aa286495.aspx) or Java technology (http://java.sun.com/blueprints/patterns/index.html, http://java.sun.com/blueprints/webservices/index.html).

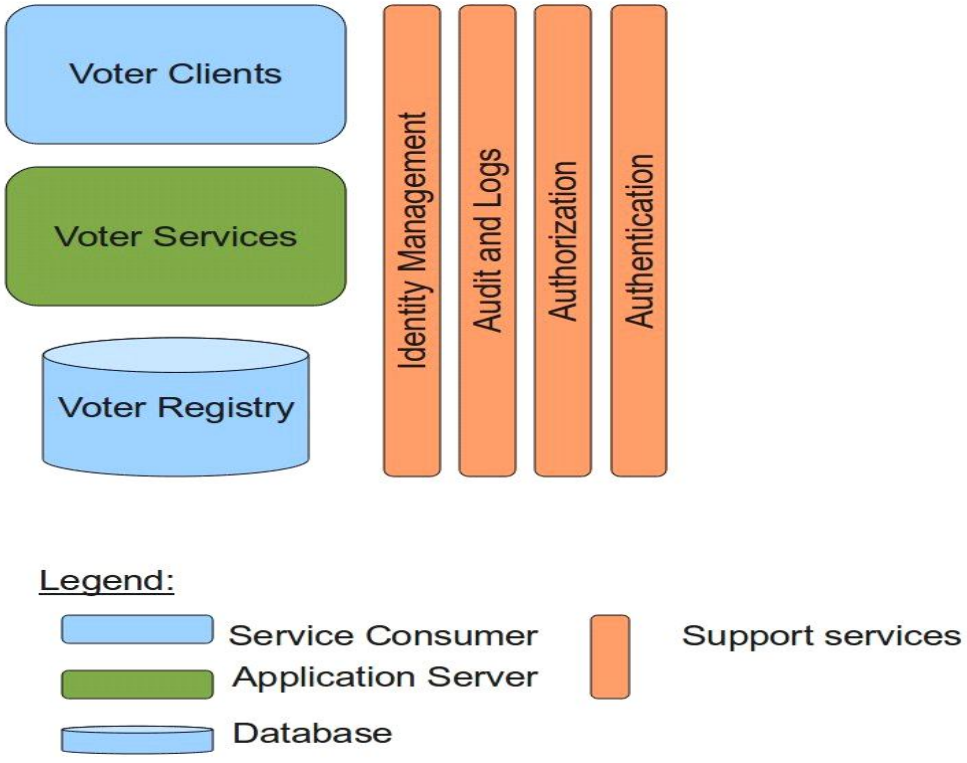

Figure 2: A three layer architecture for CVLA

Please see section 7 for detailed technical recommendations.

4.2 Recommendation 2 - External software systems

4.2.1 Summary

Recommendation: The CVLA is recommended to have an appropriate design for dealing with different external software systems.

Rationale: In addition to the roughly 75 DEO's that will produce and consume data from the central VR, it is likely there will be a data exchange with the future GIDC (Gov. Integrated Data Center). Even though the present data integration concept paper [1] envisages that other external systems interface with the GIDC, I think it prudent to design a central VR that is able to cater for later, presently unthought of external systems that would like to use the central VR data as well. This is an intermediary, quite common architectural blueprint for enhancing system maintainability.

Technical details: The technical consequences of this are that one needs a reasonable set of middleware components catering for integration. These components make the central VR able to integrate with external systems through the exposition and consumption of services as well as supporting a number of communication mechanisms like message queues, file transfer and web services.

Please see section 7 for detailed technical recommendations.

4.3 Recommendation 3 - Identity and accountability

4.3.1 Summary

Recommendation: The CVLA is recommended to have an appropriate design to deal with identity, access, accountability and audits.

Rationale: Since the VR will be a critical factor during Nepal's elections as well as having the potential of being the core basis for a future citizen registry [2], it is vital that the data in the VR can be trusted. An important part of gaining trust in a database is that it is transparent that w.r.t. accessing, modifying and deleting each data element in the registry, it is clear how that happened, who did it, what the change was and when it happened. Furthermore one must have trust in that the mechanisms are tampering proof and cannot be bypassed. The present 3 VR applications have reasonable security and accountability mechanisms built in, but are probably too primitive to deal with a general, maintainable solution catering for a more sustainable identity model and a richer set of permissions and roles. This is an intermediary, quite common architectural blueprint for enhancing system integrity, security and accountability.

Technical details: Technically, important components that underpin this, would be, at a minimum, some form of independent identity management, access and policy management, as well as facilities for comprehensive, independent logging and audit traces. There are Oasis standards for security and trust that should be adhered to.

Please see section 7 for detailed technical recommendations.

4.4 Recommendation 4 - Operational capacity for the CVLA

4.4.1 Summary

Recommendation: The agency that has the operational responsibility of the CVLA needs to be a competent, capable body that can handle ITIL type of processes in a standardized, repeatable way. My impression is that the EC IT doesn't have the necessary capability or competence to fulfill that role as it stands today. Steps are needed to make sure that either the capacity of the EC IT is strengthened or that some other organization assumes this role.

Details: Important task groups would be service, application, infrastructure and security management ensuring among other things appropriate incident and event handling, deploying applications, configuring products and managing configurations.

4.4.2 Specific recommendations

See the appendix in section 6 for a unified implications view for recommendations 4-6.

4.4.2.1 Background

The agency that has the operational responsibility of the CVLA needs to understand and be able to apply a minimum of industry accepted best practices for service management and operations. Examples of such frameworks may be the Information Technology Infrastructure Library - ITIL [8] or Microsoft Operations Framework, aka MOF [7].

An overall breakdown of roles and tasks within the MOF may be found in (cf. http://technet.microsoft.com/en-us/library/cc539249.aspx).

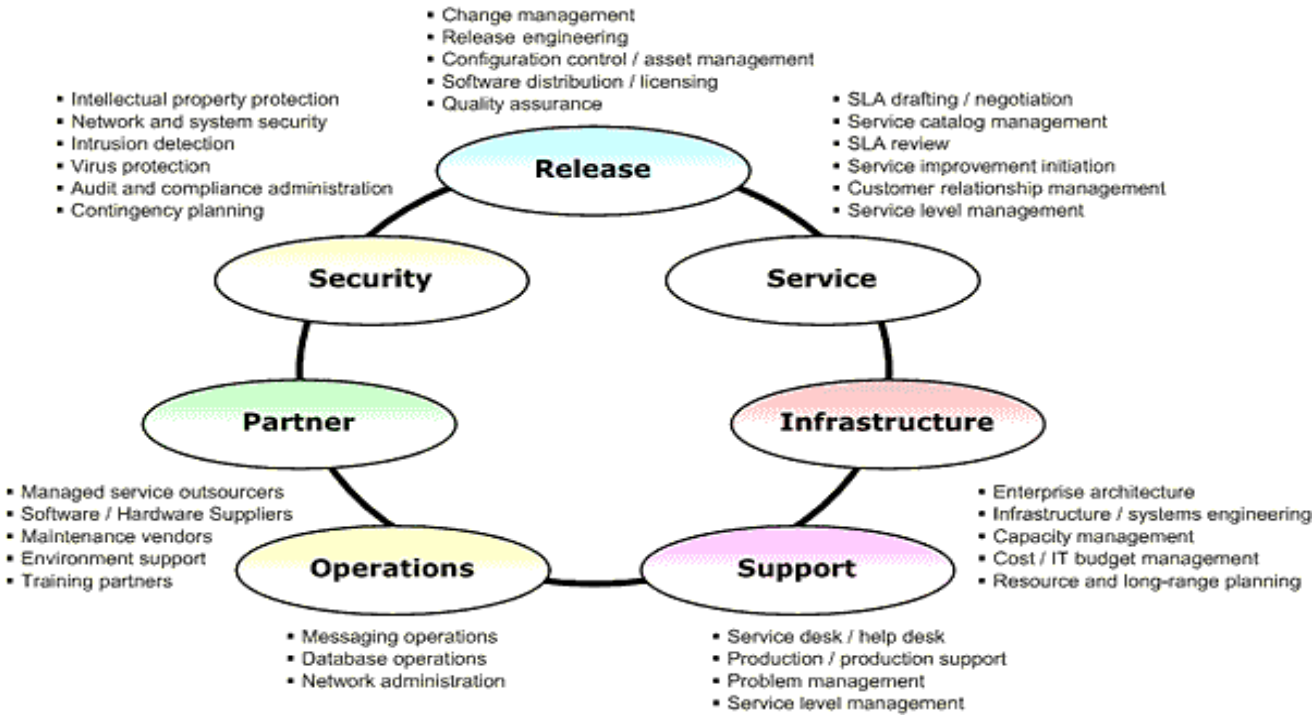

Figure 3: MOF role and task groups

4.4.2.2 A tailored solution for the EC

Taking into account the current capacity of the EC IT and private service providers in Nepal, perhaps a reasonable aggregated subset of roles and tasks to be supported are

1. A role for *Operations, Security and Infrastructure*, whose tasks would be of technical nature and pertain to infrastructure, networks, deployed applications, system administration, integration and security.
2. A role for *Support and Release*, whose tasks would be to manage application releases, migration issues, and handle incidents and events.
3. A role for *Partner and Service*, that would cater for overseeing that the services provided are of sufficient quality to the users of the registry, and to cater for the third party service supplier organization, if the EC IT uses one to run operations. The tasks here would be to follow up on and relate to the service supplier and users of the registry.

(Other arrangements are of course possible. One could for example group Support and Service together as a separate role and possibly Release and Operations as a single role.)

4.4.2.3 Organizational implications

A fundamental choice is whether (i) the EC IT should be strengthened to cater for the operational role itself, or (ii) whether a third party service supplier organization should do the job on the EC IT's behalf.

In case of option (i), perhaps 2 people should map to role 1, one person to role 2 and a person to role 3.

In case of option (ii), one person at the EC IT should perhaps map to role 3, at the service supplier's side, one person should map to role 3, and perhaps 2 people should map to role 1 and one person to role 2.

Option (ii) would probably be the more professional option, because it tends to clarify requirements and separates concerns. Option (ii) would depend however, on there being a sustained funding for entertaining a service supplier and there being such suppliers in Nepal with sufficient skills.

4.4.2.4 Skill implications

The people fulfilling the suggested roles above need to have knowledge and experience in a number of fields. The table below maps skills to roles:

Role	Skill summary	Detail reference
Operations, Security and Infrastructure	Operating Systems, Corporate Networks, Systems lifecycle management, Procedure management, Security policies and profiles, Security trade-offs, Monitoring, SNMP, Naming standards, Distributed and Remote Computing Technologies.	http://technet.microsoft.com/en-us/library/cc539254.aspx, http://technet.microsoft.com/en-us/library/cc539265.aspx, http://technet.microsoft.com/en-us/library/cc526597.aspx

Support and Release	Audit procedures, maintain and oversee applications, manage releases and deployment, server and computer diagnostics, application knowledge, manage problem tracking systems,	http://technet.microsoft.com/en-us/library/cc526608.aspx, http://technet.microsoft.com/en-us/library/cc539274.aspx
Partner and Service	SLA's, OLAs, customer relationships, partner relationships, reviews, "big picture" sight	http://technet.microsoft.com/en-us/library/cc526599.aspx, http://technet.microsoft.com/en-us/library/cc539272.aspx

All the staff, in the EC IT or outside of it, that are to assume one or more of the above roles, should be trained in these skills if they're not already reasonably proficient.

Appropriate foundation level MOF/ITIL training courses exists and should be sought for all relevant staff.

4.4.2.5 Course recommendations

For the EC IT operations role (see 6.1) I'd recommend;

- For **all EC IT** staff; Some training operators offer free/low-cost "ITIL Awareness" courses. This would be a good idea for all staff.
- For **two people** who want to specialize in Operations; "ITIL v3 Foundation in IT Service Management". *This course should preferably be in exchange of some formal commitment of staying in this role with the EC IT for at least two years.*
    - For **one of these two**; after at least half a year of operational experience (and passed exam for the foundation course), a intermediate course in some service capability aspect, e.g. "Release, Control and Validation" or "Operational Support and Analysis". *This course should preferably be in exchange of some formal commitment of staying in this role with the EC IT for at least two years.*
    - For **the other of these two**; after at least half a year of operational experience (and passed exam for the foundation course), an intermediate course in some service lifecycle aspect, e.g. "Service Design" or "Service Operation". *This course should preferably be in exchange of some formal commitment of staying in this role with the EC IT for at least two years.*

See http://www.itil-officialsite.com/Qualifications/ITILV3QualificationScheme.asp for details.

As I state above, the courses do not have to be from ITIL. A similar kind of course setup for MOF would be equally applicable.

In general I'd recommend classroom and workshop training instead of e-learning/distance learning.

It is a good idea to require that exams for these courses are passed when assigning people to these roles.

If the EC IT chooses at some time to assume an operational role itself, i.e. fulfilling the role of Agency X in figure 4, then additional training needs to be in place. I'd recommend two more people doing the ITIL foundational course and that these two, after half a year or so, take some intermediate course, one in service capability and one in service lifecycle.

4.5 Recommendation 5 - Maintenance of the CVLA

4.5.1 Summary

The organizational setup for the maintenance of the CVLA, needs to be such that both the developer and the CVLA owner have experience and competence with running software development projects.

Details: This implies among other things the ability to work under a defined software development methodology, handle iterations, changes, releases, configurations, tests and acceptance appropriately and generally ensure quality of the processes and deliveries.

1. *Recommendation:* My impression is that the EC IT doesn't have the necessary capability or competence to fulfill the client role as it stands today. Steps are needed to make sure that either the capacity of the EC IT is strengthened or that some other organization assumes this role.
2. *Recommendation:* I'm more confident, but still not sure the present Vendor has the necessary capability or competence to fulfill the developer role as it stands today. Further investigation is warranted to make a reasonable recommendation for this role.

4.5.2 Specific recommendations

See the appendix in section 6 for a unified implications view for recommendations 4-6.

4.5.2.1 Background

The EC IT must have a foundational understanding of what it means to develop software. An insight into development methodology, specifications, roles, iterations, deliverables, quality and acceptance is necessary.

For the Vendor, a more complete understanding is necessary as they will be doing the actual development work. Here detailed experience and understanding of the whole development process, both from a project management and from a technological perspective is necessary.

4.5.2.2 Skill implications

The people fulfilling the suggested roles above need to have knowledge and experience in a number of fields. The table below maps role to skills:

Role	**Skills**
Vendor maintenance - doing development	Software development methodology - in particular agile and iterative methodologies, lean development, requirements specifications, acceptability criteria, contractual handling and coordination, organizational configurations for software development, software deliverables, software

	quality, software metrics, acceptance of deliveries, non-functional and functional testing, web services, n-tier applications and their blueprints, dependency analysis, coupling, cohesion, continuous integration, automated build systems, standards for distributed applications including central oasis and w3c standards, clustering, o/r mapping, authentication, authorization, code signing, pki, service registeries, ldap/a-d, application servers, databases including logical and physical design, caching for persistence and clustering, service oriented architecture, service bus.
EC IT maintenance - keeping track	Software development methodology - in particular agile and iterative methodologies, requirements specifications, acceptability criteria, contractual handling and coordination, organizational configurations for software development, software deliverables, software quality, software metrics, acceptance of deliveries, non-functional and functional testing, web services, n-tier applications, standards for distributed applications including central oasis and w3c standards, clustering, authentication, authorization.

4.5.2.3 Course recommendations

For the EC IT Maintenance role (see 6.1) I'd recommend;

- For **all EC IT** staff; some training operators offer free/low-cost "Agile Awareness" courses. This would be a good idea for all staff.
- For **two people** who want to specialize in a product owner role; "Certified Scrum Product Owner Course". *This course should preferably be in exchange of some formal commitment of staying in this role with the EC IT for at least two years.*

See http://www.scrumalliance.org/pages/scrum_certification for details.

It doesn't have to be Scrum. Any course setup which gives good insight and hands-on experience with modern software development from a product owner's perspective will do.

In general I'd recommend classroom and workshop training instead of e-learning/distance learning.

It is a good idea to require that exams for these courses are passed when assigning people to these roles.

4.6 Recommendation 6 - IT Governance of the CVLA

4.6.1 Summary

Recommendation: Finally, the CVLA is recommended to be have a sustained, well anchored, competent IT governance body with decision-making authority. My impression is that the EC IT doesn't have the necessary capability or competence to fulfill that role as it stands today. Steps are needed to make sure that either the capacity of the EC IT is strengthened or that some other organization assumes this role.

Rationale: The likely lifespan of the CVLA is to be in years, if not decades. Whatever good decisions are being made today can easily be unmade later when other pressing concerns and less competent technical advisors are present, and the history of decisions and reasons for why things were done the way there were are forgotten. Thus, it is important that the operations and maintenance of the central VR there is a governance body that can, on a technical and business level, assure changes are done according to agreed procedure and technical constraints.

4.6.2 Specific Recommendations

See the appendix in section 6 for a unified implications view for recommendations 4-6.

4.6.2.1 Background

Proper governance for an IT project is increasingly recognized as a critical component in achieving a sustained, relevant IT solution. Built on IT Management frameworks, IT governance has grown into a separate discipline underpinned by frameworks like CoBIT [9] and recently being itself governed by a standard, the ISO 38500. Major topics on IT governance comprise how IT supports business needs, resource management, performance and risk.

IT architecture governance is a subtopic which focuses on how to manage and control software and hardware architecture.

4.6.2.2 A tailored solution for the EC

Particular areas that are important are under the control of a CVLA governance body should include

1. Changes to the business environment and capability, including organization, competence requirements, budget and expenditure, alignment with business goals (EC goals and partner (like the NID) goals), contracts and formal agreements.
2. Changes in procedures, plans, methodology, quality constraints and checklists
3. Changes in system requirements, be they functional or non-functional (e.g. response times, availability, security, accountability, data integrity/quality)
4. Changes to overall design and architecture in software and hardware
5. Deliveries, including acceptance, quality and contractual conformance.

Some of these areas are technical, whereas some are managerial. Thus, I would recommend two roles to fulfill this function;

- CVLA Technical governance for areas 2, 3, 4 and 5.
- CVLA Project governance for areas 1, 2, 3 and 5.

These would probably map to different physical persons. Note that the roles deal with overlapping areas and that the two roles would be in tight collaboration on many of the issues being worked on.

4.6.2.3 Skill implications

The people fulfilling the suggested roles above need to have knowledge and experience in a number of fields. The table below maps role to skills:

Role	Skills
CVLA Technical governance	Requirements management, issue tracking, application architecture, information architecture, business architecture, technical architecture including infrastructure and deployment, n-tier applications, web services, application servers, enterprise databases, web servers, software testing, software metrics, software quality management including response times, availability, security, accountability, data integrity/quality, foundational understanding of operations, foundational understanding of software development methodology
CVLA Project governance	Project management, agile development, contracts for software development and legal framework, offshoring, business processes, organizational capacity planning, negotiations, budgets, business-technology alignment, project metrics, project scheduling, quality systems, acceptance of deliveries management, release management, change management

4.6.2.4 Course recommendations

For the EC IT Governance role (see 6.1) I'd recommend;

- For **all EC IT** staff; Some training operators offer free/low-cost "CoBIT Awareness" courses. This would be a good idea for all staff.
- For **two people** who want to specialize in IT Governance; "COBIT Foundation Course". *This course should preferably be in exchange of some formal commitment of staying in this role with the EC IT for at least two years.*
    - For **one of these two**; after at least half a year of IT Governance experience (and passed exam for the foundation course), an implementation course would be a good idea; "Implementing IT Governance Using COBIT". *This course should preferably be in exchange of some formal commitment of staying in this role with the EC IT for at least two years.*

See http://www.isaca.org/Education/COBIT-Education/Pages/COBIT-Training.aspx for details.

In general I'd recommend classroom and workshop training instead of e-learning/distance learning.

It is a good idea to require that exams for these courses are passed when assigning people to these roles.

4.7 Recommendation 7 - Data and Application integrity

Enhanced handling of some functional areas and procedures for repeated, systematic quality control are needed for the 3 VR applications. Principal areas which should be reviewed and then possibly acted upon include;

- **Accountability features.** Security and audit features need to be strengthened.
  - Presently, audit trails are thrown away when imported from CPA to DVLA. Possibly from DVLA to CVLA as well. There is a subset of the information model (the database) that is under auditing procedures and possibly only a subset of operations that are audited. There is a practice of using a shared user, admin or superuser when acquiring data. All these undermine the general trustworthiness of the central registry.
  - Additional security features for the applications need to be put in place. In particular, we'd recommend scrambling and preferably signing the application runtimes so that they are less prone to reverse engineering and tampering.

- **Quality of entries.** There are a number of issues that needs to be looked at such that one may rest assured the quality of the entries are sufficiently high.
  - A review and test of how one handles entry lifecycle operations. For continued registration, people will be moving, there will be updates to the multimedia and I would recommend exploring the need for and form of a history/versions of records. Defined and anchored business rules defining where what operations are to be performed and what the semantics of these operations are to be, are important are in place.
  - There is a need to control the quality of fingerprint signatures/templates. Several roundtrips between representative DCA acquired fingerprint samples with particular settings and centralized deduplication and searches need to be performed. The fingerprint data needs to be realistic, w.r.t. which samples are drawn and things like grease and wear and tear of the equipment. This would help in determining a reasonable acquisition configuration (thresholds and the like) and set an expectation ceiling on the quality one could reasonably hope to attain. I would recommend setting up a test lab centrally at EC IT to do a number of calibration roundtrips.
  - The procedure for acquiring signature images make it likely the quality of signatures are at best of varying quality and possibly of generally insufficient quality. As for the above, representative signature samples must be drawn and a review of the current procedure is recommended in light of the samples' quality.
  - There is a practice of using another person's fingers when fingerprint acquisition doesn't work/is unacceptable and then writing a comment about it. One could hypothesize that another pattern would be to use the same finger 4 times. I would recommend working out robust procedures for handling these cases, establishing how widespread the practice is and possibly amending the DCA functionality to be more accommodating.
  - General text data quality control needs to be performed at regular intervals. Representative samples need to be drawn from all text fields, be validated against another

source (the forms I guess) and if there are concerns the data on average are of insufficient quality, corrective steps need to be taken.

- **Configuration and release management.** There is presently scarcely any configuration or release management at the EC IT. Over time, different versions of the application and accompanying database schema will be available and possibly running at the local, district and central levels. The deployment of new releases may not happen simultaneously.
  - A repository of all the released versions needs to be maintained centrally and available to all relevant ECN staff.
  - A defined procedure for receiving, testing and accepting new versions of the software and schemas is required such that it is clear what has been delivered and with what quality.
  - Both database schemas and applications must be versioned. They must be subject to a revision control system and they must be clearly labelled with their version.
  - A register which maps hosts (laptops, workstations) to application versions is needed. A plan for rollout of new versions on which hosts is needed.
  - There is likely to come situations were the migration of data from an earlier database version to a newer is less than trivial. The vendor needs, in collaboration with the EC IT, to work out and document how to do this procedure between relevant earlier versions and a new version of the application.
  - The structure and quality of the temporary storage of binary dump files centrally at EC IT needs to be reviewed. The structure must be such that it is easily available to all relevant personell, that it is easy to understand and that it can easily be accommodating new post-election binary dump files. Version information must here also be accompanying the data files.
  - Ease of accessibility and ease of understanding must be taken into account for all the above mentioned registries and procedures such that they can be maintained in case key EC IT staff disappears.

5 Reference

5.1 Dictionary

Term	Explanation (and reference if applicable)
AFIS	Automated Fingerprint Identification System. Presently a product called MegaMatcher from Neurotechnology.
Client	Technical term for a runtime that depends on a server.
CRUD	Create, Read, Update and Delete, cf. http://en.wikipedia.org/wiki/Create,_read,_update_and_delete
CVLA	Central Voter List Application [1]
CoBIT	A framework for IT Governance and control, stands for "Control Objectives for Information and related Technology" [9]
CR	Citizen Registry [2], synonymous in this paper with GIDC
DCA	Data Collection Application [1]
DEO	District Electoral Office
DMZ	Demilitarized Zone, http://en.wikipedia.org/wiki/DMZ_%28computing%29
DVLA	District Voter List Applicaton [1]
EC	Election Commission of Nepal
EC IT	Election Commission IT section
GIDC	Gov. Integrated Data Center [2]
GIS	Geographical Information system
ITIL	Information Technology Infrastructure Library [8].
MOF	Microsoft Operations Framework [7].
MQ	Message Queue, a particularly reliable computer communication mechanism. Many implementations and products are available. Cf. http://en.wikipedia.org/wiki/Message_queue
NID	National Identity Card [2]
SAML	Security Assertion Markup Language [14]
SRS	Software Requirements Specification [1]
the three VR applications	DCA, DVLA and CVLA
Vendor	The company doing the software development work - Wordlink/wlinktech.com
VR	Voter Registry
WSDL	Web services description language [15]
WS-Federation	Web services specification to broker identity attributes across security realms [14]
WS-I	Web services interoperability oganization [16]
WS-Security	Basic Web Services Security specification [14]
XACML	Extensible Access Control Markup Language [14]
XSD	Extensible Schema Definition language [15]

5.2 Bibliography

[1] Lindqvist P., «Software Requirements Specification for ECN Voter List Application», July 17, 2009, UNDP ESP, Nepal
[2] «Integrated Voter List, Civil Registration and National ID Concept for Nepal», September 2009, UNDP Nepal
[3] «Database Evaluation for Voter Registration System», September 5, 2010, UNDP Nepal
[4] «Implementation Strategy for Voter Registration», February 17, 2010, UNDP ESP, Nepal
[5] «Continuous Voter Registration Strategy for Nepal» (draft), August 25, 2010, UNDP Nepal
[6] «Preparation of Voter List with Photograph», July 16, 2009, UNDP Nepal
[7] "Microsoft Operations Framework 4.0", July 28, 2010, http://technet.microsoft.com/en-us/library/cc506049.aspx
[8] "IT Service Management - ITIL", http://www.itil-officialsite.com/home/home.asp
[9] "Control Objectives for Information and related Technology", http://www.isaca.org/Knowledge-Center/COBIT/Pages/Overview.aspx
[10] "The Unified Modeling Language", UML 2.0, www.uml.org
[11] "The Open Group Architecture Framework", TOGAF 9, http://www.opengroup.org/togaf/
[12] "Contract for Consultant Services for Analysis, Design, Development and Implementation Support of Biometrics Based Voters Registration System Between Election Commission of Nepal and Worldlink Technologies Pvt.Ltd.", January 28, 2010.
[13] 12 Enterprise Architecture draft documents for the Government of Nepal, all under the heading "Government Enterprise Architecture", including "Overall Enterprise Architecture document" and "Nepal e-Government Interoperability Framework (NeGIF) - Draft Report", December 15th, 2010, PriceWaterHouseCoopers India
[14] The Oasis group, http://www.oasis-open.org/specs/
[15] W3C, www.w3.org
[16] Web services interoperability oganization, Basic Profile 1.2, http://ws-i.org/Profiles/BasicProfile-1.2-2010-11-09.html

5.3 People

A list of contributors internal to the UNDP/ECN structures. Bear over with me, I might have mispelled names, and I might have misunderstood what are forenames and surnames.

Name	Affiliation	Contact details
Ram G. Aryal	EC IT	aryal_rg@yahoo.com
Sudip Aryal	UNDP ESP/EC IT	sudip.aryal@undp.org
Steve Canham	UNDP Consultant	stevecan@loxinfo.co.th
Anil K. Dutta	EC IT	anil.kumar.dutta@gmail.com
Umesh U. Gairapiplee	EC IT	umesh@election.gov.np
Sachin Karmacharya	UNDP ESP	Sachin.Karmacharya@undp.org
Prabhat Kumar	UNDP ESP	Prabhat.Kumar@undp.org
Peter Lindqvist	UNDP Consultant	peter@voxion.net
Luis Martinez-Betanzos	UNDP ESP	luis.martinez-betanzos@undp.org
Sharika	EC IT	sharika2060@yahoo.com
Matrika Shrestha	EC IT	matrika.shrestha@gmail.com
Srijana Shrestha	UNDP ESP	srijana.shrestha@undp.org
Kuikel Sudipa	EC IT	kuikel_sudipa@hotmail.com
Kiran Thapa	UNDP ESP/EC IT	Kiran.Thapa@undp.org
Oliver Vick	UNDP Consultant	oliver.vick@gmail.com

5.4 *Public document versions*

Version	Comment
0.2	Preliminary edition, sent to Luis
0.5	Sent to Peter (cc Steve and Luis)
1	Sent to Luis

6. Appendix A: Joint organizational and planning implications

From the recommendations, one may join the individual recommendation's implications so as to have a unified view on what the recommendations mean for the EC IT and it's partner as a whole.

6.1 Organization

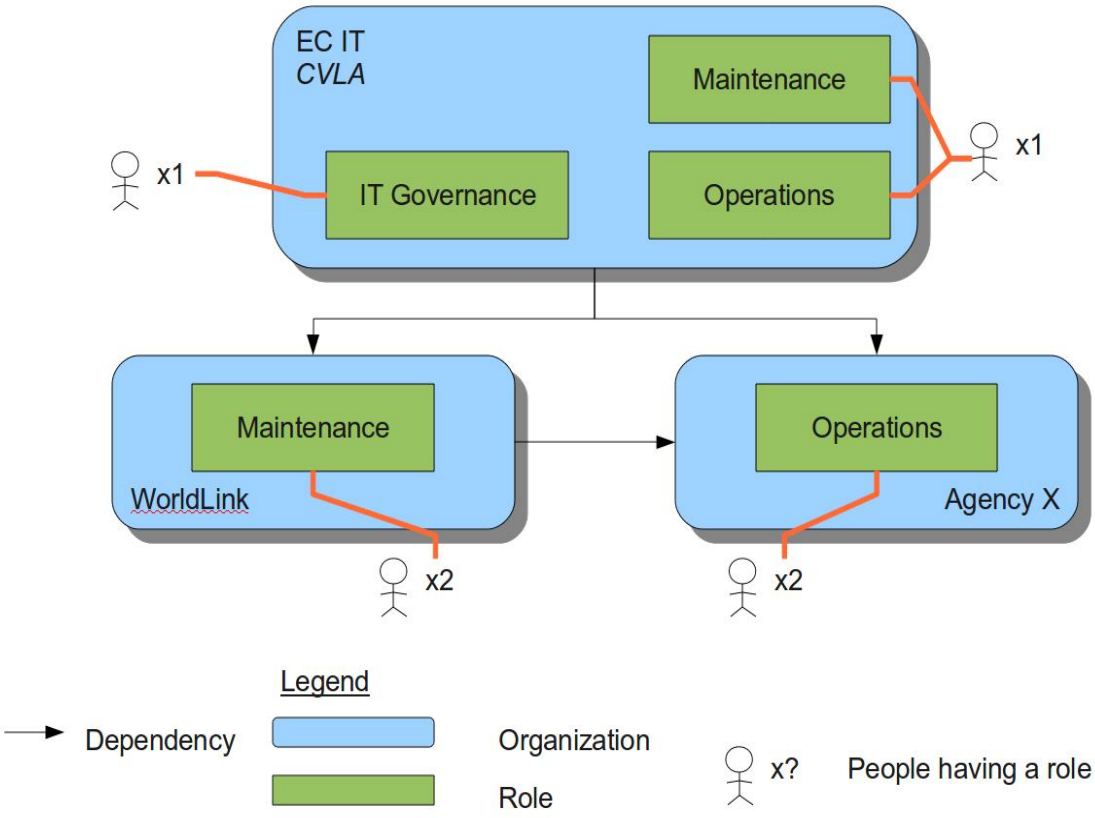

Figure 4: Suggested organization for the CVLA

There are essentially three parties and 3 types of roles.

There is the ECN, an agency doing the maintenance and development, and an agency doing operations. Then there is an operational, a governance and a maintenance/development role.

It is possible for the one agency to assume both or all three roles. In the case of EC IT as it stands today w.r.t. capacity, I wouldn't advice the EC IT to assume more than absolutely necessary. The minimum would be to do IT governance and let other agencies do operations and maintenance.

The ECN needs to work with whomever does maintenance and whomever does operations, so they need at least one person to be proficient at IT Governance, one person to be proficient with operations and similarly with Maintenance. See figure 4 for an illustration.

The agency that maintenance need at least two people being able to do the development work and communicate with the EC IT. Likewise, the agency that will do operations needs to have two people being proficient at operations and communication with the EC IT.

6.1.1 Staff recruitment requirements for the ECN

In addition to the existing work at the EC IT, we need

- at least one person to assume a joint operational/maintenance role. That is a role to *follow up and collaborate* with the agency doing operations and maintenance, not to do the work itself.
- at least one person to assume an IT governance role. That is a role to do the actual governance work and communicate with the operations and maintenance work.

Given the present capacity of the EC IT and that civil servants are likely to move one after some years of service, my recommendation would be to hire specialized long-term staff to cater for these roles. Their terms would have to be such that they would be unlikely to move on, either internally or externally.

If the EC IT were to assume the operational role itself, instead of Agency X in figure 4 then two more people would be recommended.

6.2 Plan

There is already an ongoing voter registration and consolidation process. The technology for this is the baseline applications described in section 7.1

In the following I present a skeleton plan for making the CVLA technologically and capacity wise long-term sustainable. The idea is to make a target CVLA in parallell with the ongoing registration activities. The target CVLA would be compliant with the architecture described in section 7.2.

As can be seen from the activities in figure 5 (iteration 8 and milestone 4), in time the aim is to phase out the baseline CVLA registration software and rely on a target CVLA. The scope of the plan would be to get a maintable and robust CVLA in production that can cater for a set of DVLA's, CVLA operators and an external agency like the NID.

Please note that the scope doesn't include the day to day operations of the EC IT or ECN w.r.t. the ongoing voter registration and consolidation.

6.2.1 Activities

Please see the schedule headings in figure 5 for an activity sketch.

Schedule

An estimate on the data consolidation for the VR phases

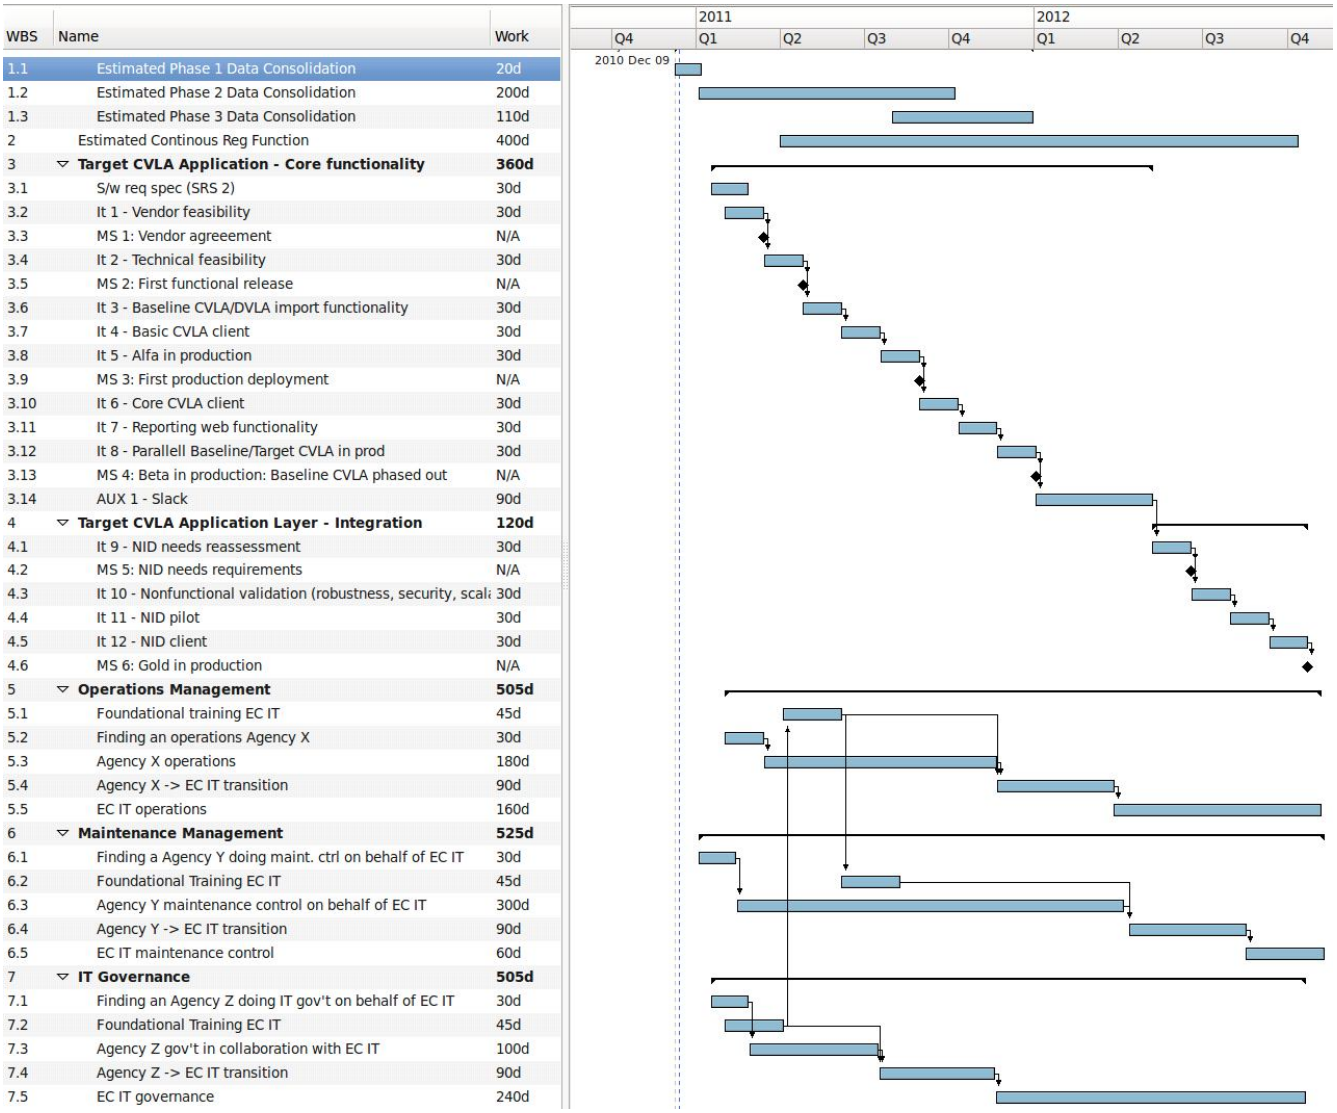

WBS	Name	Work
1.1	Estimated Phase 1 Data Consolidation	20d
1.2	Estimated Phase 2 Data Consolidation	200d
1.3	Estimated Phase 3 Data Consolidation	110d
2	Estimated Continous Reg Function	400d
3	**Target CVLA Application - Core functionality**	**360d**
3.1	S/w req spec (SRS 2)	30d
3.2	It 1 - Vendor feasibility	30d
3.3	MS 1: Vendor agreeement	N/A
3.4	It 2 - Technical feasibility	30d
3.5	MS 2: First functional release	N/A
3.6	It 3 - Baseline CVLA/DVLA import functionality	30d
3.7	It 4 - Basic CVLA client	30d
3.8	It 5 - Alfa in production	30d
3.9	MS 3: First production deployment	N/A
3.10	It 6 - Core CVLA client	30d
3.11	It 7 - Reporting web functionality	30d
3.12	It 8 - Parallell Baseline/Target CVLA in prod	30d
3.13	MS 4: Beta in production: Baseline CVLA phased out	N/A
3.14	AUX 1 - Slack	90d
4	**Target CVLA Application Layer - Integration**	**120d**
4.1	It 9 - NID needs reassessment	30d
4.2	MS 5: NID needs requirements	N/A
4.3	It 10 - Nonfunctional validation (robustness, security, scal	30d
4.4	It 11 - NID pilot	30d
4.5	It 12 - NID client	30d
4.6	MS 6: Gold in production	N/A
5	**Operations Management**	**505d**
5.1	Foundational training EC IT	45d
5.2	Finding an operations Agency X	30d
5.3	Agency X operations	180d
5.4	Agency X -> EC IT transition	90d
5.5	EC IT operations	160d
6	**Maintenance Management**	**525d**
6.1	Finding a Agency Y doing maint. ctrl on behalf of EC IT	30d
6.2	Foundational Training EC IT	45d
6.3	Agency Y maintenance control on behalf of EC IT	300d
6.4	Agency Y -> EC IT transition	90d
6.5	EC IT maintenance control	60d
7	**IT Governance**	**505d**
7.1	Finding an Agency Z doing IT gov't on behalf of EC IT	30d
7.2	Foundational Training EC IT	45d
7.3	Agency Z gov't in collaboration with EC IT	100d
7.4	Agency Z -> EC IT transition	90d
7.5	EC IT governance	240d

Figure 5: Major CVLA activities and milestones

7 Appendix B - A design sketch for a new CVLA

All illustrations and much of the terminology in this appendix is of a more technical nature. Terminology and figures are in general UML 2.0 [10] and TOGAF 9 [11] compliant.

In section 7.1 we outline how the software built by the Vendor works and how they have designed the application architecture as of December 2010. This will be the architectural *baseline*.

In section 7.2 we outline how we'd like the software architecture to look by December 2011. This will be the architectural *target*.

Note that the word "client" is used in the text in the more general sense of a runtime that depends on a server. So, a .net application with a graphical user interface reying on a database is a client, and a system being a server itself (like the NID or a GIS) would also be a client to the CVLA if it depends on it, as well.

7.1 Baseline December 2010

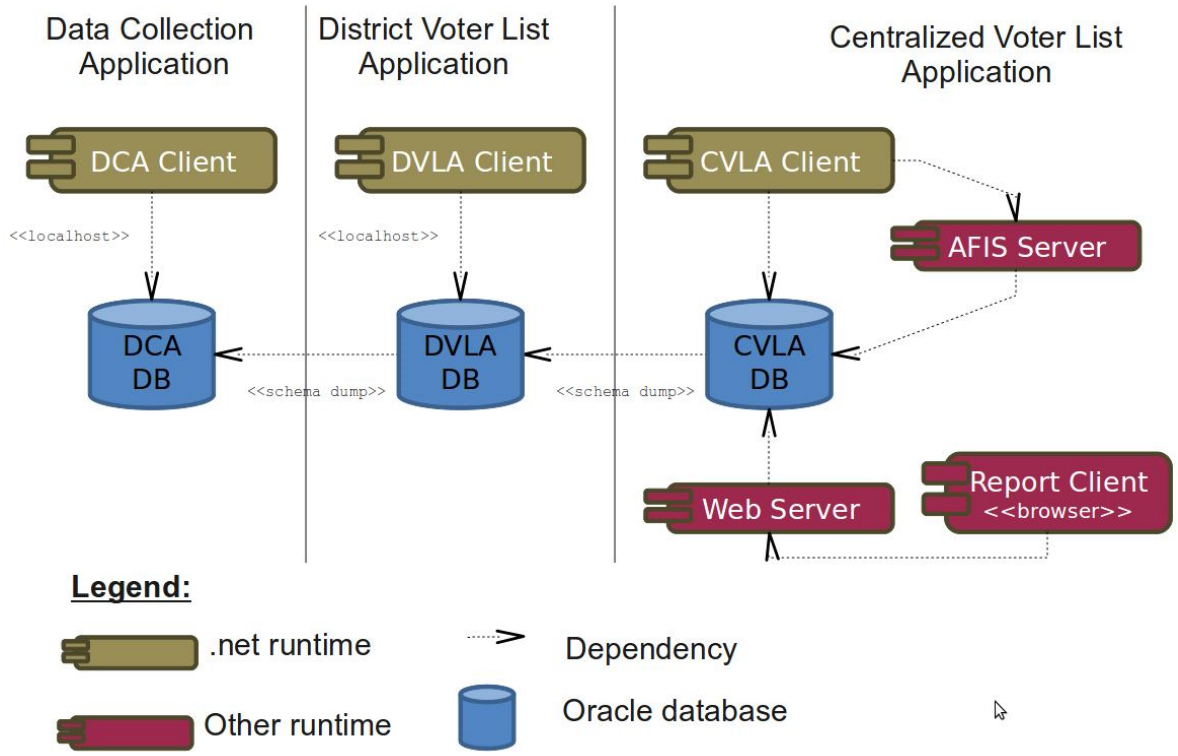

Figure 6: How it conceptually was in December 2010

There were three types of clients and databases in December 2010, the DCA, the DVLA and the CVLA. Data migrated from the DCA->DVLA and from the DVLA->CVLA by way of importing a binary schema dump from the database.

As mentioned before (section 3.3), the entire CVLA had not, in December 2010, been delivered to EC IT from the Vendor, but is based on the demo and architectural documentation showed to me at the Vendor's premises.

I'm guessing the Vendor has made all three .net clients.

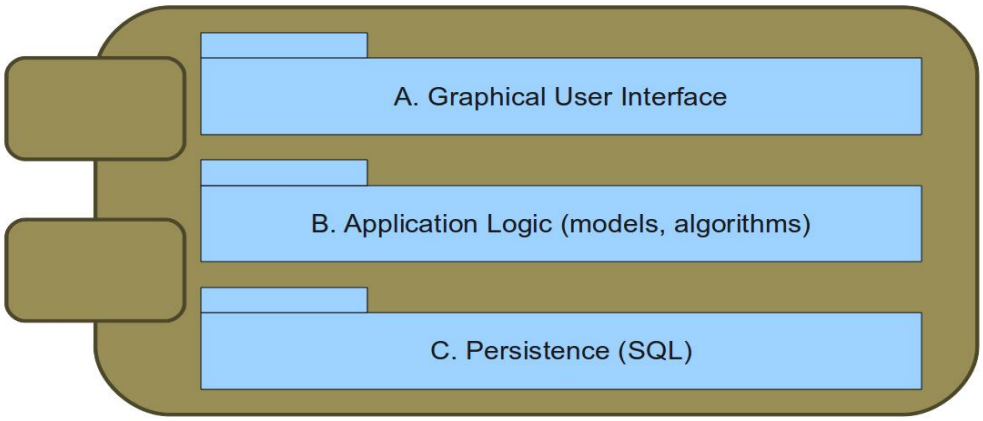

Figure 7: Logical components inside the baseline clients

For later reference I schematically non-exhaustively categorize the functional layers inside each client to include:

A. A group of Graphical User Interface functions, that at least depends on B.

B. A group of Application Logic functions, that at least depends on C.

C. A group of Persistence functions, that interfaces to and depends on the local database.

7.2 Target for December 2011

7.2.1 Application domain

7.2.1.1 General architecture

A bit depending on the security architecture, one would normally have dependencies from all clients to the Identity Service for authentication. This has been omitted from the figure for reasons of readability.

There is no reason in principle why the Identity Client cannot be a thin client (a browser) relying on the Web Server for its services. This would be an alternative and viable design.

Generally, services provided by the CVLA service and the Idenity service would be web services. Interfaces would be published as WSDL's and data types as XSDs.

Clients could be implemented using a different technology than the service providers. Thus, it is important that published services comply with WS-I Basic Profile v1.2 for interoperability.

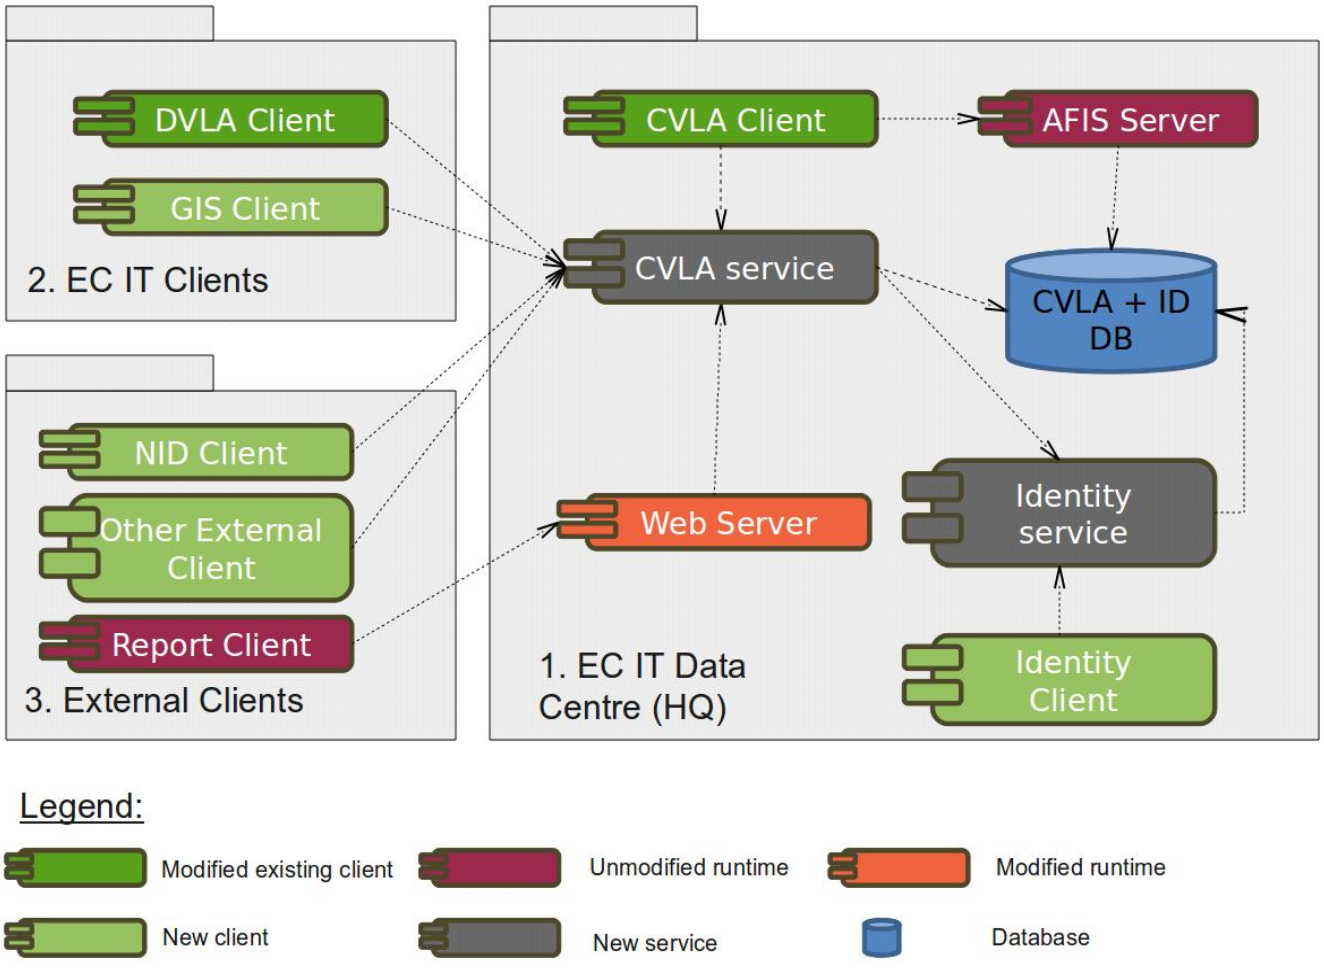

Figure 8: Target Application Architecture for the CVLA

7.2.1.2 *Security*

Generally, I recommend an architecture based on well established standards such that later extensions in functionality would be easier and less dependent on proprietary solutions.

Core functional aspects that have to covered are authentication, authorization, encryption, monitoring, logging and auditing.

Authentication, signing and encryption for web service calls should be handled by WS-Security. I'd advice using SAML as an authentication protocol, since it is a specification that supports SSO well and works with XACML for authotization (SAML v2 and XACML v2) and the standard is quite mature. The SAML specification presupposes an Identity Provider and a Service Provider. In figure 8 the Identity Provider would be the Identity Service component and the Service Provider would be both the CVLA service and the Identity Service.

In the future, it is probable that the CVLA would receive requests from external systems outside the ECN security domain. Then the need for federation will arise. SAML is quite well positioned to cater for this need. An alternative would be WS-Federation.

Every web service (and possibly other services, like ftp or message queues) published by the CVLA should be regarded as a Policy Enforcement Point (PEP) whose responsibility it is to check that the user needing services is authenticated and that the user should be granted access to the resources being requested. Thus, in figure 8 you would have a PEP in front of the CVLA service and the Identity Service. Normally, an application server would facilitate the configuration of this cross-cutting function.

7.2.1.3 Baseline client refactoring

Generally speaking, we'd aim for retaining and using as many of the already implemented features as possible in the baseline DVLA and CVLA clients.

What needs to change however is the Persistence mechanism in the clients. See logical component 3 in figure 7. Instead of the client doing CRUD operations directly on the database the clients will speak through the CVLA service component in figure 8 (a part of the application server, figure 13). The CVLA service component will then deal with the database.

7.2.1.4 The new service components

The CVLA service component will provide whatever services are needed to make all known clients satisfied. The CVLA will make available particular API's to cater for each and every client dependent on it. Examples of the services you may find below in section 7.2.1.5.

7.2.1.5 Scenarios and behavior examples

7.2.1.5.1 Scenario 1: Import a database dump of DVLA records

There are two cases. In case (i) the DVLA is offline and a dump is made available at the EC IT Data Center by physically transporting it there. In the other case (ii) the DVLA is online and a transfer of the dump is performed.

Case (i):

Step	User action	Example CVLA services used
1	The CVLA client uploads the binary dump file to an appropriate host (e.g. the DMZ server in figure 13). Transfer the dump by (s)ftp.	`requestDumpLocation():URL /* mainly for authentication/accountability purposes */`
2	Make known there is a new dump that needs to be handled.	`dumpTransferAccomplished(URL, Checksum):Status` `signDump(URL):Status`
3	An operator makes the CVLA client pick up the file and get it imported into a temporary schema at the CVLA database.	`importDump(URL):DumpRef` `getDumpVoters(DumpRef):List<VoterRef>`

4	Field and fingerprint validation is done, exercising existing CVLA client deduplication functions. The CVLA client proofing handles updates, deletes (just a type of update) and new records.	`findDuplicationsByField(VoterFields, Filter):Set<VoterRef>`
		`findDuplicationByFingerprint(VoterFPrints, Filter):Set<VoterRef>`
		`getVoter(VoterRef):Voter`
		`setValidVoter(DumpRef, VoterRef, exists:boolean): Status`

5	The CVLA client initiates a transfer of all voters that have been found valid, either new or updated ones.	`updateVoter(VoterRef, DumpRef): Status` `addVoter(VoterRef, DumpRef):Status`

Case (ii):

Exactly like case (i) but the DVLA will do step 1 and 2 above instead of the CVLA.

7.2.1.5.2 Scenario 2: Add a new Voter

The client here could be any that is authorized to do a record update. In the scenario we use an on-line DVLA, but in a future scenario [5], a VDC or some other local non-ECN agency could be requesting an update instead.

Step	User action	Example CVLA services used
1	The DVLA queries the CVLA on whether a voter already exists.	`voterExists(Voter):VoterMatch /* a set of potential matches, with a confidence measure */`
2	It seems the voter doesn't exist. Submit new Voter	`addVoter(Voter): Status`

7.2.1.6 Implementation technology and products

The application layer with its services needs to be implemented in a technology that has sufficiently wide scope, support and maturity.

There are essentially two platform choices in that respect. The application services can be realized on a (i) Microsoft .NET plattform or they can be realized using (ii) Java technology on a product/framework stack from a set of well established vendors or as open source. This report's recommendations can equally well be realized on either of these and we make no particular recommendation.

We'd advice the ECN to constrain the choices more in terms of external and internal available competence, experience and budget than any technological deliberation.

Information domain

There are three independent sources of information the CVLA application relies on. They may presently map to a single database, but are recommended to remain entirely separate. They could for performance or security reasons later be decided to reside on entirely different databases.

See figure 9 for an illustration.

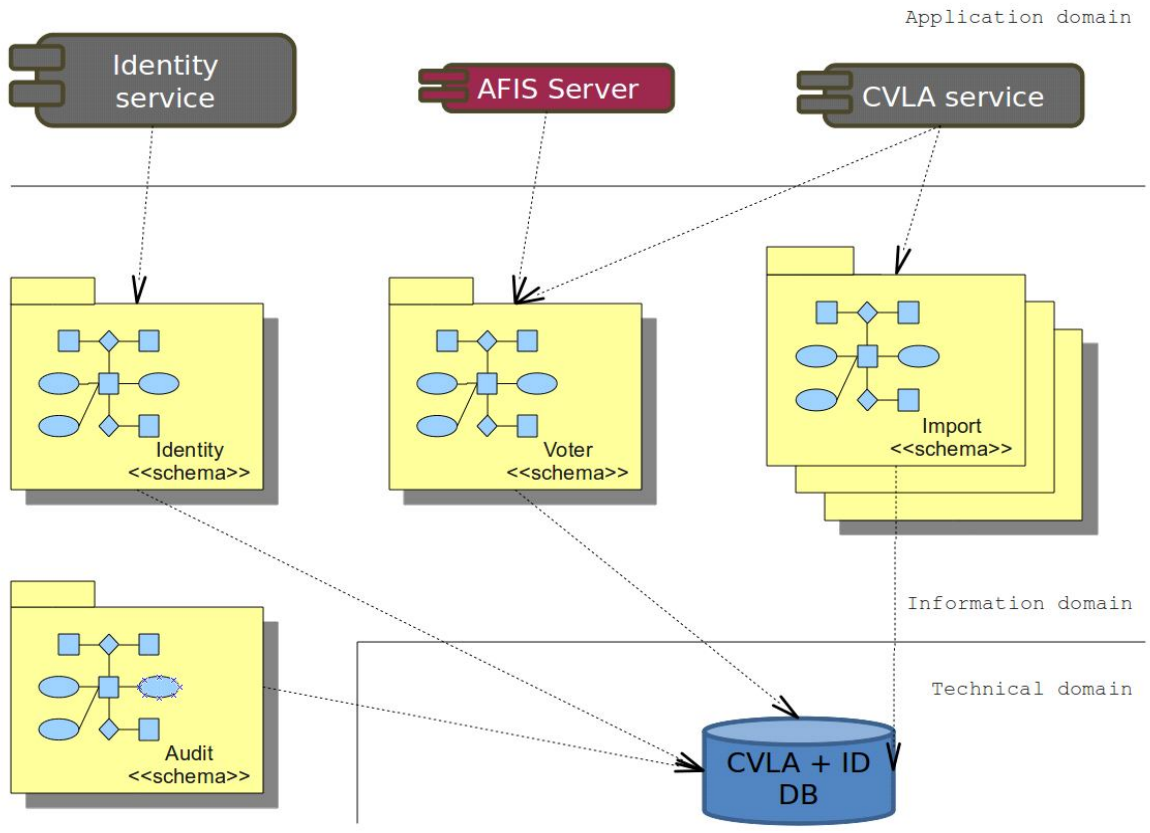

Figure 9: The four independent information models of CVLA

There is

- The *Identity information model*. This model encapsulates everything that has to do with the users of the CVLA, ECN staff, their roles and permissions. It may at some later stage serve to provision an LDAP/AD source if needed by the ECN and it may later also serve as the information store underpinning an identity federation service if that need should arise.
- The *Voter information model*. This is the core information source of the CVLA and is the Voter Registry as such. Voters, their details, all multimedia content and various housekeeping attributes are found here. Furthermore, this model allows for voter history to be captured. Thus it features a versioning of voter entries.
- The *Import information model*. This model is used for temporary storage of imported voters from a DVLA registry. There are in fact likely to be several versions of this schema, each one mirroring the schema of a particular DVLA version. See section 7.2.1.5.1 for a usage example.
- The *Audit information model*. This model is used to store traces of all or a subset of application server operations exercised and database operations. The scope and type of operations under audit should be configurable as a policy.

7.2.2.1 The identity information model

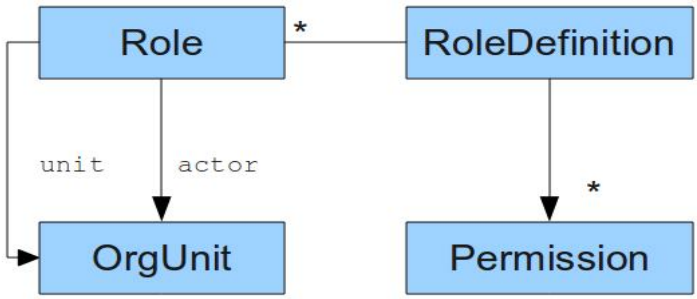

Figure 10: The core of the CVLA identity model

There is a choice between (i) using an established directory model like LDAP or Active Directory as a primary organizational information source, or (ii) to maintain a separate model and then make standard directory models rely on this. The advantage of the former is the simplicity of the approach. The advantage of the latter is the modeling freedom in the information model. One may employ a richer and more targeted model with the latter approach.

I recommend for the CVLA to use the latter approach. See figure for some suggested core entities.

A user is modeled as an OrgUnit. There is a set of RoleDefinitions, signifying the various roles users can have. A particular user would be associated with a Role instance of that definition. The user would be the Role actor and the user's organizational belonging would be the Role unit.

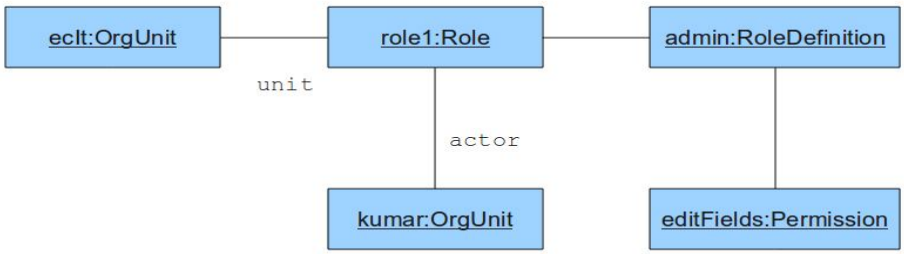

Figure 11: Kumar is an admin and belongs to EC IT

Se figure 11 for an example, where role1 makes kumar an admin with the permission to edit a field, and makes him belong to the EC IT.

7.2.2.2 Implementation technology and products

This report makes no particular recommendation w.r.t. choices of database. I do concur with the TCO analysis in [3] as there are no technical reasons why a PostgreSQL/EnterpriseDB database shouldn't work well with the recommendations in this report.

7.2.3 Technical domain

7.2.3.1 Physical architecture

In [3] we find a figure depicting a target physical infrastructure for the CVLA. See figure 12 for a copy. My interpretation is that the "vServers" are AFIS servers, and that the "Users" are where the CVLA domain tasks are performed (see section 4.3 in [1] for a list). Perhaps the "Management Server" and "Workstations" are workstations where you maintain the database and the AFIS servers.

The webserver that the Vendor has introduced for reporting purposes is possibly not in the figure.

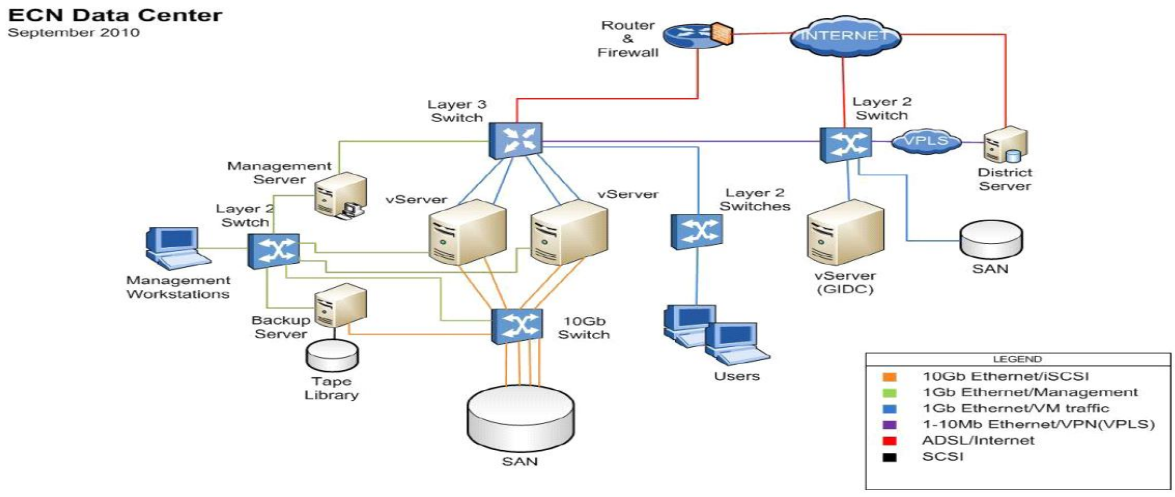

Figure12: Physical infrastructure for the CVLA as planned in Sept. 2010. From [3].

7.2.3.2 *Deployment architecture*

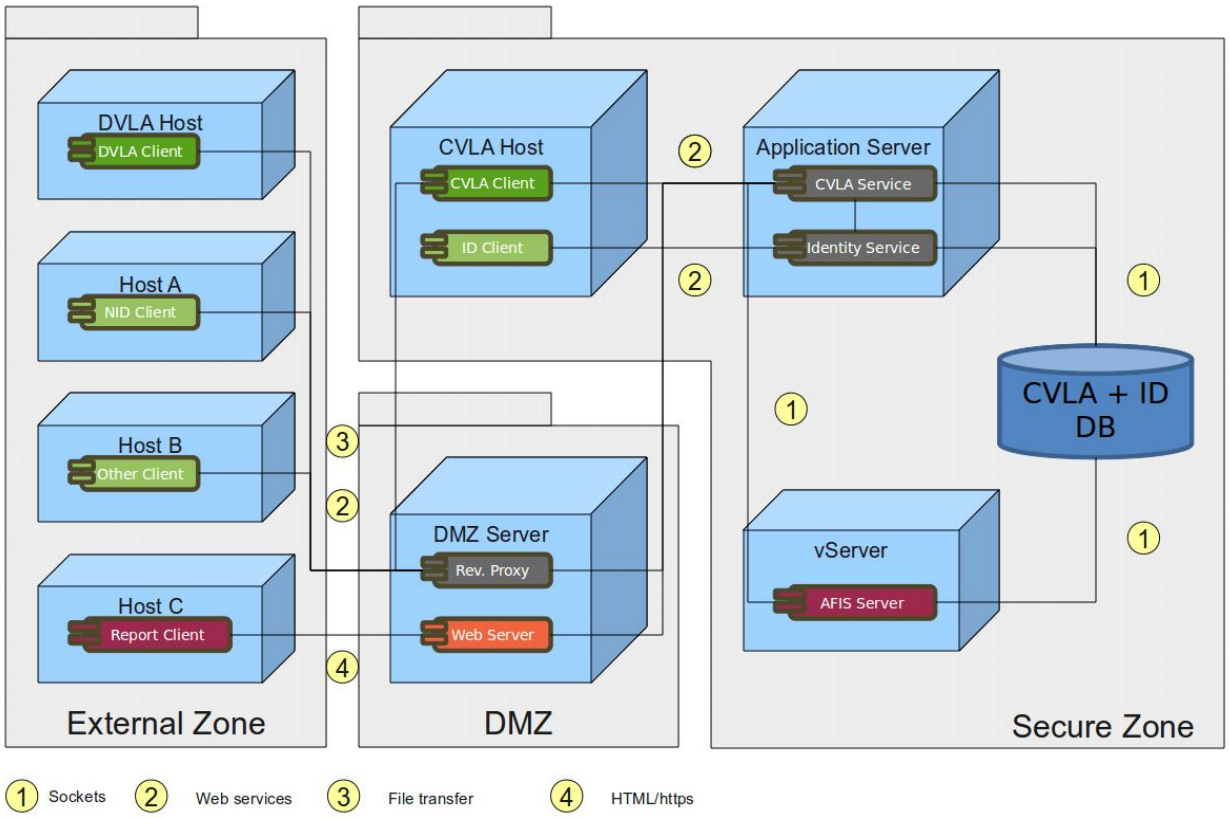

Figure13: Logical Deployment Architecture for the CVLA

The Reversed Proxy simply serves to break the incoming transport protocol, provide access control and forward requests to the Application Server.

7.2.3.3 Scenarios and behavior exemplifications

The API's published by the Application Server will in general be web services employing SOAP as a binding. Inside the secure zone and between the DMZ and the secure zone one might use MQ as a transport mechanism. From the external zone and to the DMZ it is probably a better choice to use https as transport.