All.Net Analyst Report and Newsletter

Welcome to our Analyst Report and Newsletter

Demystifying control architecture

Control architecture is, in essence, the missing link between risk management and technical controls. While many discussions of risk management are ongoing in widely read forums on a daily basis, control architecture is rarely discussed, even though it's almost always present. It appears as informal decisions that form cohesive bonds that drive implementation.

> "Control architecture may be the most complex thing to understand about enterprise information protection because it is so ephemeral and yet so critical. Control architecture goes directly to how the enterprise thinks about and acts on information protection issues. It may seem like a list of standard concepts from an introductory computer security text, but it really forms the foundations of the field, and the field continues to be rocked by the fact that these foundations are not as well understood or solid in today's environment as most people in the field assume them to be.
>
> The control architecture is typically comprised of protection objectives, an access control model, functional units, perimeters, access mechanisms, a trust model, and change controls."[1]

I'm confused – that's not a definition

Indeed it is not. As a field, information protection doesn't fully understand control architecture. It is not a widely recognized collection of concept or models, and yet it clearly goes to the issue of controls and is architectural in that it represents decisions of kind rather than amount. A notional definition might be:

> "A collection of decisions delineating models used as the basis for technical and implementation decisions."

Not really clarifying, is it? I'll put it another way. There are different aspects of protection used as a basis for making detailed protection decisions.

- If a cohesive approach is not used, a protection program will have a smattering of diverse bases for decision-making that are hard to track and manage. As a result, many different approaches to protection are likely to be used. This is not likely to scale well, will tend to produce incompatibilities resulting in exploitable gaps associated with diverse and mismatched protective measures put in place, will create difficulties in understanding why things were done and how to change them over time, and will introduce unnecessary complexity into the entire protection program.
- If a well-defined and thought out basis for such approaches are fused into a set of cohesive decisions about when to apply which approach, a model of protection will be formed that can be used to make decisions across diverse situations with consistency. The model will be more readily analyzable in terms of how it all fits together, and when making detailed decisions, all of the relevant elements of the model will be available to guide implementation. Changes wil lbe more readily handled and understood as well.

1 F. Cohen, "Enterprise Information Protection", http://all.net/ → Books ISBN 1-878109-43-X

It seems like a lot of extra effort and an aggregation of decision-making risk

Therein lies the rub. If the set of alternatives for decisions to be made are not understood in advance, or it the basis for these decisions are not well thought out, it may take a lot of time and effort to make such decisions, and as and if they change with time, the changes may ripple through to the rest of the program. Hard and fast rules without understanding and the ability to adapt leads to compliance rather than effective protection.

Most businesses start small and grow, or not, from there. Startups rarely think through such issues, and likely that is a good thing. Until a certain size and maturity is reached, a protection program is largely wasting time thinking through these issues in great depth. The cost of change is relatively low, there are all sorts of other things that can cause failures, and there is usually a great deal of risk acceptance. Just get it working to the point where we can use it, and we will fix it once we have the resources to do so. This makes good business sense in almost all cases.

But as organizations grow and mature, they start to become systematic. Somewhere they move from initial maturity to repeatable processes, and from there they may get to a defined or even managed level of maturity. As they move from repeatable to defined and managed, they often find that the complexity of protection for what they built becomes overwhelming, excessively dependent on individuals, and not readily scalable. Decisions made rapidly for startup purposes and not codified or thought through in a larger context and vision become distant memories. They are often left as legacy and not understood or detailed by those who come to manage and operate the systems later on. You may hear questions like “Why did they do that?” and “What were they thinking?” and comments like “You've got to be kidding!”.

This is the natural way of things, and there is nothing wrong with it. But as some size, typically when going from a small business to medium business, the problems start to overwhelm the benefits of not changing, and an architectural approach to protection becomes a practical way of doing things. That's the point where many practitioners fail to recognize the need for a control architecture, and end up paying the price later on.

How about an example?

I figured I would get around to that at some point... Choosing from above, let's look at a trust model. This one is particularly daunting for many because it is so poorly understood as an area of study, even though it has been studies in some level of substance.

In our standards of practice[2] we notionally codify trust in terms of sentences of the form:

> {Businesses, Content, People, Systems} x based on {historical behavior, transitive trust chains, systematic background checks, psychological factors, external clearances, contracts, nationality, group membership, investigations, credentials, certifications, size, etc.} are trusted for {purposes}.

At a high level, this is turned into less generic versions, such as these examples:

- Businesses, based on contracts, certifications, historical behavior, and size are trusted for storing backups.
- Content, based on group membership of source, is trusted for building prototypes.

2 http://all.net/ → Protection → Standards of practice

Hopefully, you get the idea. As an entity, we make the decision to codify purposes for which we are willing to rely on different bases to trust different entities. We can then further specify that as more detailing becomes necessary, until we get to an execution level, tracing it back to the control architecture decision about trust. Here is an example drill-down:

- Businesses, based on contracts, certifications, historical behavior, and size are trusted for storing backups.
  - Trusted financial partners and businesses specializing in backup storage, based on contracts meeting our standard terms and conditions and all system-specific contractual requirements, having government certifications meeting all applicable standards for the content at issue, with a history of at least 20 years without major customer-impacting incidents, and at least $10B in annual revenues are trusted for storing backups in repositories with at least the level of protection afforded by our own facilities.
    - Our bank, based on our safe deposit box contract, their certifications as an institution under Federal law, their history as a bank and in exchanges with our company, and their size ($900B in deposits), are trusted to store one copy of local backups from small facilities in safe deposit boxes.
    - Iron Mountain, based on our contract with them, their certifications [details], their history as a repository for backups, and their size ($20B a year in revenues), is trusted to store company-wide backups.

If these are codified and tracked as business decisions, they can lead to processes, procedures, and technical implementations of all sorts. The basis for decision making is laid out at increasing levels of depth through management decisions, and those decisions are codified and tracked as part of record-keeping associated with operations, including linkage back to the decisions used as a basis. When a change is anticipated, and on a regular basis, these decisions are revisited.

> If a local manager wants to make a change, they can refer to the related business decisions, track them back to the basis at the next higher level of the decision hierarchy, and eventually to the broadest level control architecture decision.
>
> If a decision is made at a higher level to change the requirements to add cloud-based storage providers with content encrypted and remove financial partners, this can be checked against higher level decision-makers and lower-level decisions for cost and consequence, and if still desired, rolled out over time.
>
> If a top-level decision to add a requirement of nationality to ownership of such facilities is made, it can be checked against existing instances for likely cost and roll-out times, and a transition made without missing any of the instances.

Now consider what would happen without such a control architecture in place. Without a basis for the trust-related decision, individuals and managers might make their own decisions and implement them on an ad-hoc and undocumented way. There would be no linkage back to the trust model, and no linkage from the trust model to the affected decisions. Change would likely be hard to manage, would likely miss instances, and would certainly be more complex. But this is only true at a scale and maturity justifying the processes and tracking.

Other areas of control architecture

We don't have a comprehensive list of these sorts of decisions for information protection. And we don't see such a list appearing any time soon. We simply don't know enough to be definitive yet. But experience and observation indicate that control architecture is very useful for, without limit:

- **Protection objectives:** For example, integrity, availability, confidentiality, use control, accountability, and transparency are selectively needed with different surety for different situations. By listing the situations (e.g., protected health information requires high confidentiality, integrity, use control, and accountability, but medium availability and limited transparency; advertisements require high integrity and accountability but no confidentiality, use control, moderate availability, and little transparency; etc.)
- **Access control:** For example, (1) clearances, classifications, and compartments are used for many government systems, (2) roles and rules may be used in many enterprises, (3) owner authorized is common across many situations, (4) subject object models are sometimes used for small scale, and (5) possession-based models are used in an amusement parks (tickets get you into rides and money gets you tickets).
- **Original identification:** For example, (1) Web browsing is often anonymous, (2) individuals may self-identify to get library cards, (3) group membership may be adequate for access during tours, (4) people may be known to others for consulting deals, (5) organizational identification cards or devices may be used for facility access, (6) credentials issued by a government may be required for airport admission, (7) tracing personal history and doing background checks may be required for hiring, (8) extensive in-depth background checks may be needed for government security clearances, and (9) forensic methods, such as DNA samples and those of parents and siblings may be required for inheritance or legal identification.
- **Access facilitation:** We start with a sentence like this one: "*{Unified, Consolidated, Independent} x {access, tracking, use control} x across {enterprises, zones, subzones, applications, mechanisms} x at {low, medium, high} granularity"* and proceed to specify in greater detail just as we do with models of trust to assure that access granted when authorized. In other words, the goal is to assure access is available when authorized.

You can imagine that there are other similar control architecture decisions. They don't directly address the protection mechanisms. For example, roles and rules may be implemented in different ways in different systems, and original identification using government issued identity cards will differ from country to country. But once the decision is made, a basis is present for rolling out a program, checking to see whether any given implementation meets the requirement, and supporting change over time.

Summary

Control architecture decisions are foundational to more detailed decisions, and as such, form a basis for making reasonable protection decisions at all levels of the enterprise. They provide the means for a cohesive program that can more readily and reliably adapt over time, and allow tracking of the basis for decisions that may be vital to controlling interactions between protective measures. They are poorly understood and yet very important to success as information environments grow and mature.

