

All.Net Analyst Report and Newsletter

Welcome to our Analyst Report and Newsletter

Cybersecurity From Scratch – Part 9: Into the weeds...

Every once in a while we get to create a cyber-security program from scratch. ...

In doing individual security inventory in preparation for implementing remote desktop, the company discovered its social media inventory... which it didn't bring up in the original Pathfinder. So... Update to scope, update to consequence analysis, update to counter-intelligence program, but since duties to protect were not yet defined, risk management was not yet completed, procedures were not yet in place, so this is just added to the build-out. Thus additions to **scope**:

Social media providers	Public relations	Positive public image, demonstration for investors and potential clients.
------------------------	------------------	---

And **Consequence analysis**:

Substantial loss of business or harm to brand if uncontrolled.	Med (consequence)	Social media content	Integrity / Availability / PR
--	-------------------	----------------------	--------------------------------------

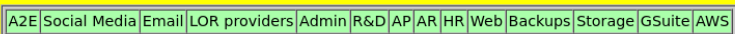
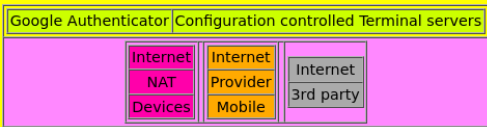
In the process, we also discovered a new person starting to gain access, and then of course had to do the inventory for their systems and uses... for the remote desktop implementation they will use.

New decisions

Of course formal approval of the previous decisions continued, but new decisions were put on pause for some development of policies, procedures, etc. to get done. This is quite common when we get to the end of the control architecture. It's a good time to take a pause and catch up on the details underlying the governance issues before pushing implementation changes.

Working through the details

As more details arise, more specific diagrams also arise. Here is an example:

Situation	Consequence
<p>Several zones for different business functions.</p>  <p>Several zones are used for different business functions</p> <p>Zones are associated with identified business functions. Users from company assets access remote desktops using Google Authenticator (2-factor for Med consequence) with access to zones based on IAM TBD.</p> <p>Cloud-based zones and temporal microzones is advised.</p>	Low Med
 <p>Cloud-based zones and temporal microzones</p>	Low Med

Architectural structures associated with consequence levels

This diagram shows the future state zones as currently known (and evolving). Users authenticate to Google Authenticator and use cloud-based temporal microzones to go from company controlled computing devices to remote desktops. From there they access cloud

service providers. The Identity and Access Management process (not yet defined in detail) uses ABAC to implement a version of roles and rules.

The roles are associated with business functions (e.g., *Social Media*) with the attributes providing finer granularity (e.g., *Twitter Admin* vs. *Twitter Content Provider* vs. whatever). A User with *Social Media* and *Twitter Content Provider* attributes as assigned by HR with approval by appropriate social media management, will be able to provide twitter content from (and only from) a properly configured remote desktop. The IAM system also dictates when/if additional authentication is required, typically based on attributes of the function and content and its consequence level. For example, Medium consequence (e.g., *Add Twitter Content*) access requires 2-factor authentication, while watching the YouTube feed requires no authentication (or use of a company computer or remote desktop). Risk aggregation then becomes a mandate of the authorization process, so that HR and social media management must agree to allow access and each has a requirement to check for risk aggregation (e.g., that no person can alter more than one social media feed). Meanwhile, change control dictates that a Medium consequence changes need to be approved, so that the suggested Twitter Content push requires approval by someone authorized (e.g., a User with the attribute *Twitter Approval*) and who is NOT the person making the change (e.g., *NOT Twitter Content*).

However, because of the small size of the company at this time, these requirements will be suspended temporarily as the company grows, and a management risk acceptance (by the CEO) will be required and revisited every 6 months to allow another process as this one develops. The decision has not yet been made, but likely the decision will be to allow users with *Twitter Content Provider* access to post to Twitter and use the intelligence process to detect the change and send it to management for review. Management review finding a problem will then deal with the issue by removing the posting and managing the personnel issue along with HR. An alternative is to have a *Twitter Content Provider* provide the content to the social media manager and allow the social media manager to post directly to Twitter with the intelligence function compensating for errors post-facto. This uses a manual process with dramatically reduced automation and complexity, but is slower. If the company were massive with a big footprint, a few seconds of bad content could see hundreds of thousands of people affected, but for most small companies, such problems occur with minimal consequence... except when the consequence ends up being more extreme because it happened just when the next investor was making their final check for investment and it put them off, or a news outlet found it and decided to enlighten the public about it.

No plan survives contact with reality (the enemy)...

Reality comes into contact with the plan, and of course, the reality changes... NOT!

The plan changes. So we started looking into the reality of options for remote desktop and controls such as those described. We started by looking at Google for the remote desktop, and of course they have the capability to do this.

After a few hours of struggling with their interfaces and mechanisms, I tried AWS. In 5 minutes I had a WorkSpaces Windows Desktop online, and was able to setup up additional ones in a few minutes each. They also have an attached file sharing solution with review and approval processes and remote access from Mac, Windows, etc. for acting like the file storage is local and automating backup of all changes. And it includes Active Directory and limited access, can be reasonably well controlled in relatively simple interfaces, has a nice

reasonably high performance remote desktop that runs on just about everything, and is inexpensive and readily expandable to hundreds of users.

So the testing began with security not particularly configured to be tight. The setup comes with Windows Defender, which meets an administrative need, regardless of how good or bad it might be at a deep technical level. It also comes with Firefox, thank goodness, because Explorer is ridiculous in every way. It has protection settings that prevent almost anything useful, it's slow as can be, and ... who cares. The point is, the Windows setup for remote desktop works well for every application so far identified in the inventory.

- This is only a test setup, so in effect, we are letting a few users try it out for performance, ease of use, compatibility, etc. Once (assuming) this passes the laugh test, we will provide a test setup to the other users for a week or so for them to be able to test things out. This is being done from my (Management Analytics) accounts rather than those of the company, and they are aware that this is only for testing purposes for now. It has no Company content, but if it works out, within a week or two we could move toward this environment for real, and deploy across Company in as little as a day, but probably less than a month, including turning on the desired security features.

Having said that, we will not be anywhere near done, because there are lots of decisions still to be made. But “crawl, walk, run” as Doug Simmons has a tendency to say. If this works out, it will be a substantial improvement almost immediately, in the sense that we will have automated backups, have some control over storage, and the ability to administer as a group. We will start to enforce the control over content areas effective immediately, at least for the higher consequence levels. For example, HR will be a restricted area of the storage solution, so only authorized users will be able to access it from only authorized virtual machines. Each user has their own VM, and of course because these are so inexpensive on a monthly basis, we can have multiple machines for each user and they can access those machines for high valued functions. For example we might have HR-configured machines for HR-authorized users who login to those machines for HR use and logout when not in use. These machines can also be restricted in various ways, for example, provided without outbound Internet access, so that while they can act on requirements, do documentation, etc. they cannot be connected to except from the remote desktop application, and they can only access pre-loaded applications and file storage content areas they are authorized to.

This it looks like the plan is coming together and will be able to function, if we can pass the current testing level. Some will ask why we didn't got to Microsoft Azure instead. The real answer is we haven't gone there yet. We have what looks like a satisficing solution. If it works, the time and effort to explore other alternatives is too expensive at this point to leap into it. Nothing we are doing depends particularly on AWS, and there are alternatives. Depending on assurance requirements, we may also duplicate it elsewhere, and we will almost certainly test it elsewhere, after we have made these improvements. This will provide a take-out capability for a commodity solution, which is ultimately a requirement for a cloud-based solution.

This decision means we have to move our secondary backup solution from AWS, likely to Google. Resiliency, reliability, performance, usability, cost, and more... all factors that have to be satisfied, which is why satisficing works a lot better in this situation than seeking optimization. We could negotiate about it forever, but by picking one trying it, we have saved time and money. This is precious in smaller companies, and all the more so in startups.

Integration and transition

One of the constant challenges in building an IT architecture for a small but hopefully rapidly growing company is the eternal challenge of integration. Each new application, content type, person, and business relationship has to be integrated into the information environment. This includes identification of consequences, and from there, risk profile and matching taking into account duties to protect, and from there any implications to management and control architecture, and then to the technical details.

As companies grow, they have staff engaged in all of these activities, and as a result, they learn to work with each other, create workflow mechanisms, and so forth to manage and operate the system. But in small and startup companies, the staff is not there, and the complexity is only starting to arise. Until they form a systematic approach, everything is done for the first time, and with limited resources, they cannot afford to become systematic in the big picture right away.

The approach we have found most successful is to do something the first time and get it right once. Then document is at a light level, having made lots of mistakes along the way. The next step is to repeat the process and document as you do it. After 3-5 times, it becomes repeatable with documentation in place, which leads to Defined maturity. That is the goal for this company over 6 months, and should be the goal of most startups developing a cyber security program with at least Medium consequence levels over their first 6 months of operating their cybersecurity program.

Over time, design patterns emerge. These design patterns become standard approaches that are integrated into the operations to make them more efficient and effective while limiting their granularity, but also the related costs. This develops into economies of scale over time. This is similar to the way roles and rules work, and part of why roles and rules emerged as a good approach to addressing protection. The new version of this, ABAC (attribute based access control) replaces roles with attributes for a finer grain of control mapped into the multiple roles and other things that emerge from usage.

The first step is low granularity control, because we don't know what the finer granularity will look like until we get there in this particular company. This is driven by the evolution of what they do, how they do it, and what management decisions are made with respect to the control architecture and implementation. Perhaps more importantly, the mechanisms put in place become the trained activities of staff, and thus the integration of training and awareness creates barriers to change. As companies grow, the number of people using the same methods creates a sort of culture, and culture is even harder to change because it is part of group cohesion, both in terms of things they have in common to complain about and things they find in common in other ways. Group cohesion is part of loyalty, which is an important part of defense against insider turning behaviors, and ultimately disgruntled employees. This then goes into the lifecycle program and HR-related issues.

A very good reason not to go wild with protection changes is because of its disruptive effect on people and operations. This leads to more mistakes, misunderstandings, stress, and reduction in performance, both by individuals and through synergistic effects, the company as a whole. Of course if you are looking to change a culture, massive changes are one way to affect that change, at the risk of all the disruption it brings.

Changes yet to come

It is becoming clear that there is additional content in the company that has not been covered in the inventory yet. The current inventory is at a high level, and this means that as we move forward it will be augmented in small ways. But there are also major functions not yet identified (or at least put in the inventory). One was identified earlier in the week, but we are awaiting the details.

The transition to remote desktop will also provide a far more complete inventory both at the level of the minutia of what's actually used, and at the larger granularity of business functions, as access controls and collaboration start to come into play. It also provides the potential for mechanisms to track it far more effectively (and efficiently). Here's what we expect will happen:

- Testing will get to the point where people have confidence in the solution as workable.
 - They will say YES for the next steps.
- A transition to the cloud will take place for select areas of the business and all users will have virtual desktop access for those functions.
 - The users will have some really minor transition issues, but will also identify things they cannot do. This will translate into added details in the inventory, including needed software packages, files they didn't realize they were using, etc. As these are added, the protective issues associated with them will appear and be resolved.
- Storage will be transitioned to the cloud for everything users are aware of on their systems, using remote file shares from their local systems.
 - At this point, there will be inventory items that fit in known categories, and inventory items that are not yet associated with a category. Because of the access controls over these things, there will be a slowly decreasing number of these sorts of items over time as new categories are added and transitions made.
 - Along the way, we will discover that there are many copies of the same files in different systems, and the file storage solution will start to resolve these at a project level. The details will be worked out over time, but this will also support the retention and disposition process as the association of people, content, systems, and businesses becomes clarified by the organization of things across the company. This will help resolve legal holds, backup and restoration processes, and record keeping overall.
- Remaining mechanisms will be transitioned to the virtual machines and local systems will become largely access points and temporary areas for working offline.
- Users will be able to transition to purely company-owned resources, or perhaps a decision will be made to use company-owned VMs on private systems, or even private VMs for personal use on company-owned assets with controls to assure that the users' privacy is maintained and they can take out their content as/if desired.

Conclusions

Things are starting to move with the temporary authorization of bypassing separation of duties till an IT director can be hired at the company. Defined maturity remains feasible in 6 months.