Begin Note Box
This chapter was written by Jeff Rochlin, who spent many years working in an organization that did work for the DOD and strictly followed Military Standard 480B (MIL-STD-480B) regarding what they call “configuration control.” I first met Jeff while consulting for that company twenty years ago. I learned so much about change control working with Jeff in this environment that I could think of no one better to write this chapter.
End Note Box
The purpose of this chapter is to explain the business and political aspects of the process of building or enhancing your data protection system. Equally important to designing a good technical solution is getting buy-in from everyone that will either benefit from or pay for it. For more on the technical aspects of designing or refining your data protection system, see Chapter 16.

Data protection is not the sexy part of IT. It reminds the organization they are vulnerable to various risks that often have nothing to do with IT’s core competency. The resources that need to be applied are costly and, in most cases, don’t show up in the final product you are selling to your customers. You are selling your organization on an insurance policy that deep down, no one wants to buy. It might not be easy, but the truth remains that your data protection plan is going to be one of the most important investments you will make in any organization.

Before you go off and spend big chunks of a budget on data protection, you need to make sure what you build properly covers the needs of the organization. After all, there is no sense in having an insurance policy if it won’t cover your losses in the event the unthinkable happens. So, in these next pages, let’s put structure to a process and develop the tools to build an effective plan.

What Does your Organization Do?

To do this right, it’s not going to be enough to be a strong technologist with vast knowledge of what goes into a data protection system. You need to understand the purpose of your organization, as well as any external requirements that may be placed on you via laws and regulations.

- Are you a governmental organization? If so, what services do you provide and functions do you serve that IT makes possible?
- Whether you are a governmental organization, a commercial or non-profit company, does your organization provide products and services through an e-commerce model?
- Do you work for a commercial business that produces a physical object that has to go through an R&D, Prototyping and Manufacturing cycle before it gets to market?
- Are there external requirements for data protection that may also drive your design, such as regulations or laws that require you to store (or not store) certain data types in certain ways?

Each part of this complex system is going to have differing requirements. You as the data protection person will need to architect the best solution that protects everything for everyone.

Start with an understanding of the organization and the services or products you provide in order to understand the importance of your organization's data. Find someone at the organization that is willing to take the time to answer any questions you have about what the organization actually does and what the full portfolio of products or services you are providing, as well as how they are delivered to your customers. Having the full picture will help you drive the journey you've undertaken.

Before you can begin processing the information you will gather, you will need to build yourself a small organization designed to process it, as well as a document system to record the information. This framework will be the basis of your new system, and it is the purpose of the next section.

Build Yourself a Framework

Data protection impacts all aspects of the organization. You will need input and approval from many different groups, both technical and non-technical. With that in mind, you should prepare to recruit a few teams of people and involve them in a series of review boards that will participate in providing requirements and feedback on design and operation of the service you are building.

As you meet with these groups, you will need to generate documentation that will inform the process and ultimately act as your future reference. Here are some standards you should adhere to when building these documents:

Document Templates

Start out with a template for all documentation you will be creating in this process. There are some basics in the structure of each document you create.

Purpose statement at the top
    Explain the purpose of the document as concisely as possible, in no more than one or two paragraphs.
An executive summary
    If the document is designed to provide a design or conclusions of testing, you want to make sure the people who will need to provide approval actually get the information to make their decision. We all know they won't read the entire document.
A revision history

All documents should be treated as a living document and therefore, like life, will be subject to constant change. Especially during certain active periods in the project, those changes could be happening multiple times in a week (or even a day). A small table that notes a document revision number, the date of the revision, the author of the revision and a couple of bullet points on what changed in the revision will help you keep track of what document you are looking at and how you got there.

A sign-off page

Accountability is critical when developing a program that will be as critical to your organization as data protection. You want to make sure that all the critical approvers and subject matter experts (SMEs) are willing to commit to the plan by putting their signature on the final version of the document. Having this sign-off will also signify the completion of the content (until it needs to be updated again).

Policy/scope

The policy or scope of work being addressed can also be valuable to define the specific subjects the document is meant to address.

Glossary

A glossary to clarify any terms that the document addresses is particularly helpful for the non-SMEs involved in the approvals process.

Appendices

Any other documents or supporting information that relates to but is not necessarily directly a part of the document should be attached as well to the end of the document.

Review/Advisory Boards

As I write this, I'm thinking about the two dozen or so technical editors that will review everything in this book. It's amazing how that many points of view can affect a project. A good review structure for your system will help ensure success by bringing in diverse viewpoints that prevent you from missing some critical components or requirements. There will be several iterative phases in this process, including the following.

Requirements Review

The requirements review that will include members from various departments, including a senior management sponsor to make sure there is overall approval for the project from the organization. The Chief Information Officer (CIO) is a good choice for this as they both understand technology and the strategy for how the organization uses technology.

Design Review

The design review board (DRB) will include members for the technology-specific teams that could provide insight into the way technology is implemented in the organization. (Some organizations may call this an architecture review board, or ARB. The purpose is the same.) If you are a large organization, make sure to include systems engineering, database engineering, storage and network engineering and cyber-security. They will be able to review the infrastructure design and make recommendations on how to improve the integration and operation of the new service. The design review process should

include a preliminary design review (PDR), where the plan is first reviewed against the requirements to make sure all is effectively covered and a production readiness review (PRR) when it is all built and final testing is conducted. The purpose of the PRR is to get everyone to take a last look and make sure nothing was left out.

Operations Review

You should pull together the operational teams that will run the service in production and allow them to fully understand what they are being asked to do. When the operations review is finished, you should have a runbook that will act as the user manual for the system.

Change Review

Finally, there should be a change advisory board (CAB) that exists as part of the technology organization and reviews all changes before they go into production, where they can impact the daily operation of the organization. The CAB acts as the gate-keeper for all changes to protect the integrity of the organization. After the operations review is completed, it should go before the CAB before it goes live.

Project Management

Using sound project management practices will help you coordinate the work, provide the available resources and help get the work scheduled, keep you accountable to the scheduled deliverables and help ensure a smooth roll out to production.

Bring in the project management office at the very start. These are the folks that track and gate all the work that happens in most technology organizations.

Collecting Requirements

If you work in an outfit with five people, all sitting in the same room, it's going to be easy to look across the table and ask what is really important to getting the job done, but most of you aren't living in that world. So first and foremost, it is going to be critical to identify your key stakeholders and understand what they need to function effectively.

What is RPO and RTO?

There are two critical metrics that drive any data protection plan: the recovery point objective (RPO) and the recovery time objective (RTO). Both topics are covered extensively in Chapter 4, but a short version is RTO is how quickly you need to recover operations after a disaster, and RPO is how much data you agree you can lose in the event of a disaster.

Find the Subject Matter Experts.

Take your knowledge of your organization's services or products and create a list of all your internal customers by department. (BTW, if you are in data protection, everyone in the organization is a customer.) Build a list of the SMEs who can best describe what each group

does, what they need to have protected and how critical it will be to have it on-line. They will fall into a couple of broad categories.

Data creators

Where does your data come from? What departments is it from? Is it generated by intelligent systems used in a manufacturing process best understood by operations and people on the manufacturing floor? Is it created by a team of highly skilled artists/writers/editors at a substantial hourly rate? Does it come from a sales department, customer service, or any other department that is directly facing the customer? Understanding which groups are generating the data is the first step in understanding how complicated it would be to recreate the data from scratch.

Data creators will come from your production and operations teams, product management, organizational/business intelligence and data services. You should include a representative of the compliance and cyber-security teams as well, since they also have crucial requirements on how data gets stored and used.

Remember there are multiple entry points for your data. You may have to worry about the customer, service, product, inventory, order and historical sales databases at the same time and each change at a different rate based on the internal workflow of the organization. It could become impossible to keep all the data always protected and on-line at once without taking very costly measures to protect it, so it is critical to work with the data creators to understand what the RPO looks like. That will give you the picture of what a restored system actually looks like (for example, we can’t recover to the moment of failure but maybe to failure minus one hour).

Be sure to ask questions about how many events or transactions they deal with in an hour or day so you get an understanding of the churn rate of data in the systems. Remember to collect the information on where data is stored and how much actual space it uses when talking to the database/data services teams.

Executives

You are going to need to talk to members of the executive staff because they will have the best insight into the speed your organization is operating at. Understanding timelines for expected deliverables is critical and the non-technical leadership (i.e. department heads) are the ones that can best share that.

Any self-respecting executive will tell you they want everything protected all the time, but they will be the best at helping you understand the ebb and flow of the organization, which will help you determine priorities based on your discussions with the data creators.

Be prepared. These good people will tell you that no downtime is acceptable for the organization. (That is until you present them with the bill for a fully redundant, always active

system that is geographically diverse across multiple data centers that keeps data synced to all sources in real time.) So, be prepared to have a serious discussion about what costs the organization can support, with an understanding that money is always an object in the discussion. Armed with this knowledge you will be able to determine your RTO.

Compliance and governance

Make sure that whatever you are doing with data protection also complies with any laws and regulations that pertain to your organization. Privacy has become a major issue that is being addressed by new government legislation around the world. GDPR, the legal framework implemented by the European Union has requirements that an organization must be able to completely delete a users information upon request, this would also include information stored on backups and archives. The California Consumer Privacy Act (CCPA) requires an organization to be able to report back all data that contains customer information in all systems, including backups. Seek out a Subject Matter Expert from legal or governance teams of your organization to help make sure that your design follows the rules to be able to access data as needed in backups and archives.. For example, your organization may have a data protection officer (DPO) that handles compliance to the general data protection regulations (GDPR) in the European Union. They would be a natural SME for this area.

Solicit Requirements

Armed with your list of SMEs across the organization, set up an interview with each one to get their views on what the organization's requirements are. Keep in mind that you will be talking to both technical and non-technical folks, so be sure you have someone there that can act as translator in case the specifics get too deep for your audience. SMEs are not usually generalists, but (as the title implies) experts on specific areas of the organization.

Be prepared when you go into the meetings with whatever documentation and diagrams you'll need to explain the concepts you need them to understand in order to answer your questions. Try giving them an example using their processes and then tell them a piece of it is now gone. How would they handle it? Your `S:\` drive or `OneDrive` is gone. How would that affect your operations? That gives them an easy way to wrap their head around what you are asking.

Be respectful of their time and make sure you schedule appropriately. (For some it may be better to have three or four 20-minute meetings than to have one two-hour meeting.). Never forget that while your job is critical to the long-term success of the organization, they have day jobs that are critical to the immediate success of the organization.

Meet with the groups individually so their views and requirements are not immediately influenced by any others. There will be time to hash it out at the requirements review later.

Review Requirements

One you have made the rounds and collected all the requirements that each department believes are critical to survive a data loss, you need to get everyone on the same page. You should have enough information at this point in the process to figure out where the data in your organization lives and how much of it there is. Don't be afraid to go back to the SMEs if you need additional clarification. This is really important stuff, so you want to get it right.

You should also start to have an idea of how quickly data is being generated and how quickly it is changing. (Note: The data services team told you this.)

You should have an understanding of how long services or products take to create or deliver, so you can understand the tolerance for a partial or full organization outage. (Note: The management team told you this.)

Take the copious notes you gathered and put them in a presentation, then invite the key stakeholders to review. These are going to be representatives from the management team to speak to the organization and data creators along with some key members of the technology teams that run the infrastructure, so they can speak to the data. It's best that this happens around a table so everyone can ask and answer clarifying questions.

Start out by defining the problem you are trying to solve and then lay out the requirements you heard from each of the various departments in a presentation that shows each expectations'. You are not designing the solution yet, but you are clarifying what each part of the organization believes is critical to protect in order to keep things running in the event of a disaster.

PowerPoint is your friend, but don't get too flashy. Keep the slides high level and be prepared to spend more time talking to them and not asking others to read and come to their own conclusions. Remember, this is about verifying their requirements and showing the best blending of it all into something unified. Try very hard to avoid having this discussion turn into a solution session. Too many chefs spoil the soup.

While it isn't necessary to have a budget and quotes at this point in the process, it is a good idea to have some ballpark idea of what things will cost so you can help your audience understand how requirements translate into cost.

Service level agreements

This is a good point to lay out the service level agreements (SLAs) you will be establishing in order to meet the RPOs and RTOs you agreed to. Remember that data protection will include heavy use of network resources, storage devices (solid state and otherwise) and possibly even tape. These are all resources that have some form of physical constraint that will cause your

service to take time and money. Also keep in mind that data protection typically utilizes more network bandwidth than any other service in your organization.

For example, if you copy your data to the public cloud to protect it, it must move across a wide-area network (WAN) that will have some bandwidth restraints. Once there, the cloud can become expensive if you use a lot of it, so you may want to figure out a point in time to move data to a slower tier or even delete it. When you need to restore it to meet your RTO, you have to balance the amount of time it will take to physically restore the data as well as what data will need to be discarded for congruency, against a higher cost to reach your goals. Be sure to do your homework and define what that means to the service level you are guaranteeing for your customers, so you set their expectations properly.

Talk about a charge-back model

A *charge-back model* means that a given department will be held financially responsible for the amount of the service they use. You should introduce this as a topic of discussion during the requirements review. It will impact the design of your solution.

For example, the marketing department may generate hundreds of gigabytes of user data in the course of their work and don't typically take the time to clean up files they don't need after processing. If the requirements expect that all data be protected, knowing that they will be held accountable on the budget to pay for that space can help them decide if they really need it all protected. A charge-back model helps drive home the reality that all this infrastructure to protect the data isn't free and should lead to a discussion about data classification.

Data classification

It may become necessary to take time in this process to classify the data that is being protected. Not all data is created equal and it is likely a fair portion of that data can be thrown away without impacting the normal operation of the organization. Taking the time to complete an exercise that determines what data is critical, important, nice to have, and expendable will directly impact your RPO and RTO. Having said that, many data protection systems have been designed with a single data classification of "important," so don't be surprised if that happens to you.

But be sure to emphasize that even if it costs them money to protect their data, they shouldn't leave anything out that they really need protected just to reduce costs. Remind them that job one is about saving the organization in the event of a problem and if it isn't protected, it can't be recovered.

Wrapping up the requirements review

Check your ego at the door. Each person at the review is going to have their own opinion and will also be approaching the problem from the perspective of their part of the organization. Do not let leading questions or unfortunate comments feel like a personal attack on you or your

work. This is likely the first time they are hearing the priorities of other parts of the organization, so there should be spirited discussion to clarify any misconceptions about what is best for the organization overall. If everything is top priority, then nothing is top priority so remember, you are guiding them through a process of understanding. Leave at least ten minutes for each person you invite to the meeting for the discussion and make sure everyone understands that the conversation will determine the requirements of the organization.

Take good notes. If the consensus of your conclusions needs to be revised, be prepared to update the presentation and meet to review it again. Iterate as often as is necessary to get it right. You are setting the ground rules for how the organization will protect itself and recover in the event of a data loss, so it's important to get it right.

Once the requirements review is completed, put the conclusions into the document template and pass it around to everyone who attended the review for a physical signature. (You can do this through a system like SharePoint as well, so long as it captures an official digital signature and freezes the document). It is very important that you take this step. People will take an extra moment to make sure they understand the information in front of them when they will be held accountable for the results, and nothing says accountability like putting your signature on the dotted line.

Last thing to do with the Requirements Review Document is ask for people to participate in the DRB. Odds are you will need fewer senior management and more technology and data creator types for the design reviews. It is always valuable to have people who helped develop the requirements involved in reviewing the design.

Throughout this whole process, your role is to suggest things to others and see what they think. This is how you build consensus. You ask them questions to solicit their requirements, then you tell them what you think they said their requirements were. Once you all agree on that, your next step will be to suggest possible ways to meet those requirements. In other words, *it's time to actually try to build this thing.*

Design and Build Your System

Now that you have everyone's requirements and made sure they know that you understand them, it's time to move forward in the process and actually try to build something that meets those requirements. Don't worry, you're not going to do this alone, either. As you go through the design phase, your goal is to get their consensus on a design, just as you did when gathering requirements. This will start with having multiple ways to meet their requirements, each with different advantages and disadvantages.

Draw up Multiple Designs

You are going to draw up multiple ways to meet their requirements – different designs with different price points, different actual recovery times (RTAs) and actual recovery points (RPAs), as well as different levels of requirements for those who will use them. (Some will be easier to use, and others not so much.) Your job is to draw up these designs and guide the interested parties to a consensus around one choice.

Your first plan should always be a “pie in the sky, money is no object” solution that achieves the RPO and RTO defined in the requirements. This is your blueprint and the best example to show what a perfect solution would actually cost.

Then value engineer it into a second-best solution that still meets the objectives but has some built in caveats. Reduce upfront costs as a trade-off to extra cost on the back end when executed.

For example, say your organization makes animated movies. Your creative teams (data creators) will generate thousands of files and run complex compositing and rendering processes against them to create thousands of final images. In the process, you will create millions of small files that go into producing those images. It may be sufficient to save only the files created by the creative team to meet your RPO, but you will need to re-create the millions of others in order to get to where you were before the data loss and land at your RTO?The computer and human time to re-process those files will add a cost, but it may ultimately be cheaper than the extra effort and cost to capture the millions of files in the initial backups.

Keep in mind that failures are rarely clean and orderly events, so the effort to clean up the restored files to a state that allows you to flip the switch and get back to pre-failure state may actually be more prohibitive than just re-processing some lost files.

Be sure to list and explain the trade-offs that come with being fiscally conservative. Be prepared to demonstrate that the cost delta between the perfect system and other options should account for the extra cost needed to clean up and reprocess data to get back to steady state.

Your time to get back to business (RTO) will have to justify the amount of data you have to restore and clean up (RPO) so be prepared to make the sales pitch. The truth always lies somewhere in the middle of two viewpoints, so be prepared to cover all the scenarios in your solution.

Review the Designs

You’ve studied the whitepapers and industry best practices. By the time you get to this point in the process you will have read Chapter 16 of this book, which goes into detail on the technical

aspects of designing or updating your backup system. You've reached out to potential vendors and sat through the dog and pony shows and received budgetary quotes. You've thought on it a little bit and finally produced a beautifully written fifty-page document complete with scenarios, diagrams, dataflows, cost analysis and a final recommendation. It's time to get some validation that the solution makes sense to everyone involved in operating and building it.

Time to call together the DRB. Remember, the DRB includes members from the technology specific teams that could provide insight into the way technology is implemented at your organization, including systems engineering, database engineering, storage and network engineering and cyber-security. You will also want to find the SMEs from the world of the data creators as well. Look to your production and operations teams for this.

Start out with a summary of the final requirements document and go deeper into the woods with this presentation by getting into the specifics of where data is being stored, how it is being encrypted, how much bandwidth you expect the system to consume, as well as how it will need to be operated. Be sure to note your expectations for the RPO and what additional work will be required to get to RTO. None of us know everything about everything, so the input you take from the technical folks will help you refine the design. Bring your sharpened pencil, because your SMEs will provide great feedback.

Take the feedback and iterate again. If you came away with a major change to the design, architect it and run it past the DRB again. If you do some sandboxing or a full proof-of-concept and things turn out to be different than expected, document it and run it past the DRB another time. When you have the final design, you are prepared to move ahead with, fully document it as per your framework and send it around for sign-off.

Select and Build the System

This is the fun part. You get to buy a whole bunch of things and make them work together. You'll configure the system to look at the data sets and process them. You'll meticulously time everything to certify that you are meeting your SLAs defined in the requirements document. You'll run it in parallel for a few weeks, or a month. Then you'll build and run a full-scale test that will prove you have achieved your RPO and RTO goals. (Remember that achieving the recovery goals is really the only reason you're going through this process.)

Congratulations. You have a data protection system. Now it's time to build your operational plan and documentation.

Document and Implement the New System

No job is complete until the paperwork is done. If you design the most beautiful data protection system, but no one but you knows how to run it, you haven't done your job. You must document the system in such a way that people who didn't design it can run it without your intervention.

Defining Operational Responsibility

Everyone must know their responsibilities in the new system. Start by making a *responsible, accountable, collaborator, informed* (RACI) chart that will delineate what teams are responsible for the different tasks associated with the operation of the service.

Responsible
: Those who do the work to complete the activity

Accountable
: Those who are held accountable for the completion of the task or deliverable

Collaborator
: The person responsible for the activity must collaborate with in order to complete the activity.

Informed
: Those who are kept up to date on the progress of the activity

You'll define the various tasks that need to be carried out by the different teams. Figure 2-1 is a typical RACI chart. It makes it very easy to see who is responsible, accountable, collaborating, and who should stay informed.

	Systems Admin	Data Ops	NOC	Head of I.T.
Run Nightly Job	R	A	C	I
Data Outage Incident Mgt	A	C	R	I
Quarterly Testing	R	A	C	I

Figure 2-1 Typical RACI chart

Defining and securing approval from the responsible teams in advance of roll-out will go a long way to making sure things go smoothly. Be sure to pull together the operations review board

(ORB) to review the RACI chart and answer any question or concerns about how your new system will impact the organization.

Operations Review / Documentation

Once you have a working system that you want to start using against your organization's critical production data, it is time to make sure all the documentation is up-to-date and all your audiences are addressed. You already have the requirements document and the design document, so it's time for the operations manual, runbook, or standard operating procedures (SOPs). These are all different names for the same thing. They allow people to understand how this fits into the day-to-day running of the organization. It's time to get an operations document in place, and as the designer of the system, you're in the best place to make that happen.

Be sure to meet with each team defined in the RACI chart, such as systems administration or the network operations center (NOC) and collect some requirements from them on what they need documented in order to take on the operational responsibility for the service. It will probably be really helpful to you and them to bring over someone from each team, train them on their responsibilities and have them take a first pass at writing down a manual. Having the perspective of the person who will incorporate the task into the day to day is invaluable.

Documentation is Good

Let's take a moment to discuss that giant elephant in the middle of the room. No one likes writing documentation. We all have more fun things to be doing at our jobs. It's just easier to manage the system yourself than to document it so others can run it.

You are going to have to be a salesman on this journey, so be prepared to explain that documentation is critical to the efficient and orderly operation of the system and a runbook/manual/SOPs written from the eye of the person doing the task will always be more effective than one written by someone who never has to do that. The following paragraph is your sales pitch to get people to write documentation.

> So, here is why the greatest thing you can do for yourself is to write your documentation…. You want to go on vacation? Take a day off work? Sleep through the night? Get promoted to a new and better job? You can't do any of those things if you are always the one they need to call when something isn't working. If you are under the impression that it gives you job security, I can assure you that every single one of us can be replaced on a moment's notice with the proper application of capital resources. So why not actually make your life easier and help build the documentation that will allow an operator on the graveyard shift to solve the problem before waking you up.

Runbooks

The operations runbook (i.e. SOPs, operations manual) should incorporate the same template as the design and requirements documents. As a matter of fact, be thorough and attach them as appendices, so a curious operator can learn more about the service. There should be a service summary, a revision history and a sign-off page for every department that will be participating in the operation of the service.

In addition, the runbook should be made up of checklists that define the regular tasks by frequency that the operations folks will carry out. They should be formatted like a checklist so that a busy operator can always grab a copy and check the boxes when they have it completed.

There should be a frequently asked questions (FAQ) section that goes in depth on any process that can get complicated.

Be sure to include a contact list that includes all the major vendor support contract information in the event of a component failure, as well as a list of the responsible members of the various groups that can be impacted by the service and the executives that need to be kept informed. You should have mobile phone information and be sure to note how they prefer to be contacted in an emergency.

Finally, there should be a section in the runbook for the operator to list any incidents that take place, with a brief summary and resolution. Leave room to refer to the tracking support ticket. This will be really helpful for an operator to see if something has happened before and how it was dealt with.

My personal opinion is that at least one copy of the runbook should be on paper in a binder where your operations folks can find it. We live in a world of cloud based services and Wikis to hold our documentation and none of them are guaranteed to be available in the event of a system or network failure when you need the runbook for restoring your services. It may feel like a waste of trees to print it, but you will feel much better about it when you are standing in the darkened machine room and need to remember the order that your servers and storage need to be re-booted to bring the services back on-line.

Implement the New System

Now that the system is designed, tested, and documented, it's time to officially make it part of the computing environment. The CAB is going to have something to say about that.

If your technology organization doesn't have a CAB that meets regularly to review changes before they go into production and potentially impact customers, they should. Having a process like this will dramatically increase overall uptime by asking a few simple questions:

- What are you changing?
- Has it been thoroughly tested?
- What services can/will be impacted by the change?
- How do we back the change out in the event something goes unexpectedly?
- When is the change scheduled to be made and how much time will it take to be implemented?

A CAB will provide visibility to anyone in the organization to any changes so that problems could be noticed faster and rolled back in a timely manner if they cause issues. If you don't have a CAB, create one. Also helpful would be a change manager whose sole responsibility is the CAB and everything it oversees.

When you are ready for production, bring your complete documentation set to the CAB, especially the runbook, and be prepared to review it all. If you've done your work well, many of your CAB members will have been on various review boards and will already understand what you are bringing on-line.

This is also an iterative process, so follow the queues of the CAB. If they have concerns and need more data, collect it and review it with them again. If they tell you to hold changes due to other changes on the slate, be patient. It's always harder to find the root cause of a problem when you make too many changes at the same time. Once blessed, go live.

Anytime you need to make a change in your service, such as a software upgrade or a restore test, make sure the CAB has been informed as well. Trust me on this, they are your best friend for organizational stability.

Takeaways

Congratulations. You have built and implemented a world class data protection service for your organization that protects the critical assets of the organization and accounts for the necessary requirements that keep the doors open. Plus, you have done it using a methodology that has produced a thorough set of documents that will be useful for the lifecycle of the service. You have made life easier for the operations crew, so running it and troubleshooting problems will be clearer and require less experience. You have proper documentation to support the needs of the cyber-security team when analyzing data anomalies and repairing damage after potential breaches, as well as satisfied the compliance groups when they need to satisfy the organization is properly supporting it's SOX or GDPR data protection and exclusion rules.
By thoroughly understanding the organization, its requirements and all the components of a good data protection solution, you will build one of the most important and often ignored components of a solid, successful technology department.

Now that you have a thorough understanding of your organization requirements, you know what needs to be backed up and what needs to be kept for very long periods of time. Now what you need to know is the different types of data protection systems you can use to do just that. This means knowing the difference between backup and archive, which is the purpose of the next chapter.