RUBRIK PRESENTS

THE GORILLA GUIDE TO...®

Backup & Recovery Best Practices

Gary L. Olsen

INSIDE THE GUIDE:

- Why Traditional Backup/DR Is Failing You
- Fighting the Scourge of Ransomware
- Stay One Step Ahead of the Bad Guys

HELPING YOU NAVIGATE
THE TECHNOLOGY JUNGLE!

ActualTech Media
www.actualtechmedia.com

In Partnership With
rubrik

THE GORILLA GUIDE TO...®

Backup & Recovery Best Practices

By Gary L. Olsen

ACTUALTECH MEDIA
6650 Rivers Ave Ste 105 #22489
North Charleston, SC 29406-4829
www.actualtechmedia.com

PUBLISHER'S ACKNOWLEDGEMENTS

EDITORIAL DIRECTOR

Keith Ward

DIRECTOR OF CONTENT DELIVERY

Wendy Hernandez

CREATIVE DIRECTOR

Olivia Thomson

SENIOR DIRECTOR OF CONTENT

Katie Mohr

PARTNER AND VP OF CONTENT

James Green

WITH SPECIAL CONTRIBUTIONS FROM RUBRIK

Kristina Avrionova, Principal Product Marketing Manager

Robert Hamilton, Sr. Director, Product Management

Adam Eckerle, Director, Technical Marketing

ABOUT THE AUTHOR

Gary L. Olsen has worked in the IT industry since 1981. He has authored two books on Active Directory and numerous technical articles for magazines and web sites, including Microsoft. Gary worked for Hewlett Packard Enterprise as System Software Engineer and Solution Architect from 1992-2017, and was named a Microsoft MVP 15 consecutive years. Gary retired from HPE in 2017 and now works as a technology consultant and freelance author.

ENTERING THE JUNGLE

CALLOUTS USED IN THIS BOOK

SCHOOL HOUSE

The Gorilla is the professorial sort that enjoys helping people learn. In the School House callout, you'll gain insight into topics that may be outside the main subject but are still important.

FOOD FOR THOUGHT

This is a special place where you can learn a bit more about ancillary topics presented in the book.

BRIGHT IDEA

When we have a great thought, we express them through a series of grunts in the Bright Idea section.

Takes you into the deep, dark depths of a particular topic.

EXECUTIVE CORNER

Discusses items of strategic interest to business leaders.

ICONS USED IN THIS BOOK

DEFINITION

Defines a word, phrase, or concept.

KNOWLEDGE CHECK

Tests your knowledge of what you've read.

PAY ATTENTION

We want to make sure you see this!

GPS

We'll help you navigate your knowledge to the right place.

WATCH OUT!

Make sure you read this so you don't make a critical error!

TIP

A helpful piece of advice based on what you've read.

INTRODUCTION

Welcome to The Gorilla Guide To...® Backup and Recovery Best Practices! At its heart, this book is about keeping your business resilient and running, even when facing more threats than ever.

Backup strategy and infrastructure is a huge challenge for IT managers in today's complex, hybrid cloud and legacy onsite environments. To take just one example, the pandemic has forced companies' workforces to work remotely, creating new challenges. There are many boundaries that have to be negotiated, such as public and private clouds, on-premises infrastructures, and geographically dispersed endpoints. An increasing amount of data governance and regulation presents additional hurdles. All these issues profoundly impact backup and recovery strategy, making IT managers consider the question "'where is my Backup and recovery strategy now, how successful is it, and where does it need to be?"

IT managers, staff, and CIOs, alike, should employ best practices in this guide to effectively ensure business continuity. An effective data recovery strategy requires an enterprise-wide, holistic view to protect and recover data from a wide variety of sources. This includes modernization, addressing security to protect the data from attackers, such as ransomware, adapting to changing workload landscape and new applications, and how implementing cloud services and automation techniques can provide powerful solutions. All of these issues will be addressed in depth in this guide.

So, let's get going! We start with an overview of the drivers that have led to the need to modernize backup and recovery.

CHAPTER 1

Modernization of the IT Infrastructure

In This Chapter:

- Moving to a Cloud Infrastructue
- The Modern Backup and Recovery Strategy
- Service Level Agreements

The Information Technology age, especially in the past 20 years, has moved forward in a meteoric rise. This is seen in software defined operations, hardware technology, applications, and virtualization, faster networks and data growing exponentially, especially with mobile devices, and all these requiring more sophisticated data collection, management, and storage. See Figure 1 for a timeline of important technological advances in the past 20 years and how they affect the enterprise.

The point here is that many companies have experienced the events in this timeline, but have failed to modernize their systems, infrastructure and strategies to keep up. Things that hinder moving to modern technologies include company size and industry, IT Staff training, budgets, and time. Looking at the acceleration of advances over the past 20 years, what will the next 20 present? Or even the next 5 years? IT managers should ask if they are prepared for new and complex business needs, such as a geographically dispersed workforce and hybrid and multicloud adoption. If you have a plan in place, can you protect it, and recover from security attacks? Having a cloud first strategy, that will support these new computing, needs will help.

2000

- Windows 2000 released. Most companies were on Windows NT
- Backups were still mostly on 9 mm magnetic tape, for mainframe and larger systems. Smaller servers utilized disk and CD/DVD storage.
- USB drives announced—storage capacity of 8MB (yes, MB)

2001

- 3G networking launched – data transfer 4x that of 2G
- VMware announced first x86 server virtualization product

2003

- Initial release of Zen – the first open source x86 hypervisor
- Microsoft Virtual PC
- Microsoft SPOT smartwatch released

2004

- …thru 2011 IBM Watson developed for AI question answering, used today for healthcare, finance, legal and retail applications

2005

- Tablets introduced
- HP releases Integrity Virtual Machines for HP-UX
- VMware releases VM Player

2006

- Amazon's Amazon Web Services announced (beginning of cloud computing)
- VMware releases VMware Server (free)
- Microsoft releases Virtual PC as a free product
- HP releases Integrity Virtual Machines v2.0 suporting Windows Server 2003

2007

- Touchscreens
- SmartPhones (Apple iPhone)
- Hitachi announces first 1TB Hard Disk

2008

- Hadoop defeats supercomputers - fastest system in the world for sorting terabytes of data
- Google processed 20 petabytes of data in one day
- VMware releases VMware Workstation 6.5 for Windows and Linux – accelerated graphics on Windows XP guests
- Google speech recognition – using parallel nueral networks, spotting patterns in huge volumes of data streaming

2009

- 4G first deployed – enables high quality video streaming
- Cloud-based Network Attached Storage
- Fitbit Tracker launched

2010

- Microsoft Azure released
- OpenStack announced

2012

- Sony SmartWatch released running Android
- Google Compute Engine
- 2013 Kingston releases first 1TB flash drive
- Google Glass released

2013

- Snapchat reports 700 million photos shared daily

2014

- Smartwatch sales total 4.2M

2015

- Google stores 10B gigabytes of data
- Apple Watch released
- Smartwatch sales total 19.4M devices

2019

- 5G launched for public use

Figure 1: This graphic shows some of the major technological advances over the last 20 years

Drivers to move to a Cloud Computing Environment

Cloud Computing, while a relatively new term is not particularly a new technology. However, as today's networks and hardware advances, it has become increasingly popular. Think of the cloud as a configuration of servers, storage, network components, and the like that can be hosted internally or externally and managed internally or by a contracting company. An IDC report, "Rubrik: The Data-Forward Enterprise—How to maximize Data Leverage for Better Business Outcomes," states that 70% of the CIOs surveyed indicated they have a "Cloud First" strategy. Let's define the cloud environments and determine what will drive organizations to adopt a Cloud Strategy. There are four primary cloud environments – Private Cloud, Public Cloud,Multicloud and Hybrid Cloud.

Private Cloud

The Private Cloud can be as simple as a single rack of equipment, such as a converged infrastructure, multiple racks of equipment, or even hosted in a remote location by a managed service. **Figure 2** is HPE's private cloud that runs on their "Synergy composable infrastructure."

Figure 2: HPE's plug-and-play private cloud solution

This is essentially a plug-and-play cloud solution. All major manufacturers have similar offerings.

The infrastructure, and usually the location, is owned by the company (likely as Capital Expense (CapEx) and managed by a contracting company or internal IT staff. There are advanced hardware solutions sold by major vendors to make this a turnkey operation. The Private Cloud is often more a term that is used than an advanced technology, however, this is a good place for a company to start in the Cloud.

Public Cloud

The public cloud is gaining in popularity in the past 10 years or so, especially since Microsoft's Azure came on the scene to compete with Amazon and Google. Unlike the Private Cloud, the Public Cloud offers a "pay as you go" model. Rather than purchasing or even leasing equipment, predicting storage needs, hiring additional staff, and leasing floor space, a service is purchased to provide infrastructure for specific needs. It could be a valid solution, for example, for a new project that needs eight Windows Servers, 10 TB of storage, and has certain Service Level Agreements (SLAs) to host mission critical apps.

Public Cloud providers typically use a "shared responsibility model" where the service provides infrastructure and the customer provides apps, security, backup and recovery software, management services and so forth.

A company such as Amazon, Google, IBM or Microsoft can provide the public cloud services in stages—such as turning on two servers with 3Tb data storage, and then turn into more in six months without incurring charges until they are operational. Note that this can include installing the operating system, applications, backup and related services, and licensing according to the customer's needs. Is there a sudden need for more than 10Tb? There is no longer a need to ask a storage admin for more space, or even another server, as you can allocate resources yourself via your public cloud account. In addition, you will not need to purchase and budget for more hardware. You can manage it yourself, have the service manage it, or both. If the project ends after 12 months, terminate it. No legacy equipment to dispose of, and no long-term licenses to pay for.

The Public Cloud is not just for large companies. In fact, it may be more beneficial to a small, startup company to put their environment on a public cloud service and simply pay the fee rather than an IT staff—at least to get started.

Multi-Cloud

With the expansion of cloud offerings, many companies are employing a multi-cloud environment, where they have multiple cloud services hosted by multiple Cloud providers. For example, Google may have a new service offering that Microsoft does not have or Amazon may give better pricing. In addition, multiple layers of resiliency can be achieved within a single cloud service or across clouds to mitigate failures in the cloud. The aforementioned IDC report stated "… we expect organizations to use multiple public cloud services, referred to as multi-cloud. Even if organizations are not currently operating in a multi-cloud environment, we expect that the vast majority will do so."

Hybrid Cloud

The Hybrid Cloud is simply a combination of Private and Public Cloud environments, including on-premises resources. Although separate entities, they are tied together providing benefits of both models. A Hybrid Cloud can also refer to a collection of multi-site resources with managed or dedicated service models and public cloud resources that could involve multiple service providers.

One of the primary contributors to adoption of a cloud—especially a public or multi-cloud environment is *Data Sprawl*.

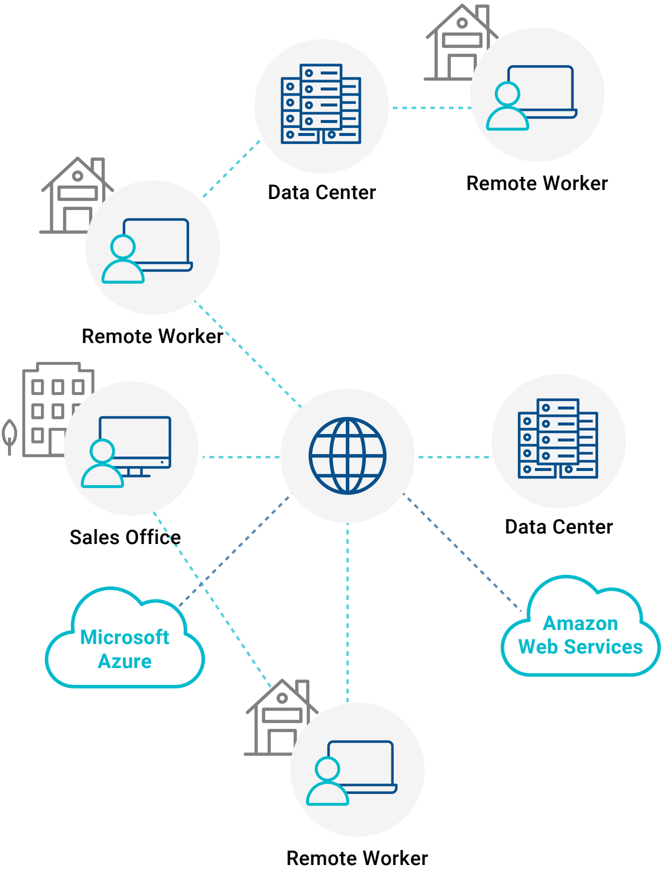

Figure 3: Data sprawl in the modern IT environment

Data Sprawl

One of the key drivers to a cloud environment for backup and recovery is data sprawl. Twenty years ago, there was little attempt to backup data that left the office, and storage was concentrated on-premises typically in individual servers or small storage arrays. In today's environment, however, as shown in **Figure 3**, the modern IT environment has data in many different locations.

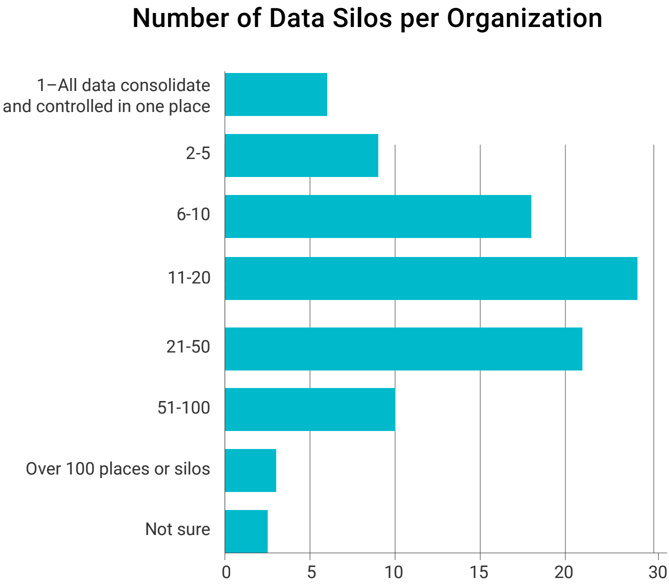

n=801 Base=all respondents Average=23

Source: IDC's Rubrik Thought Leadership Study, December 2019
Notes: The survey is managed by IDC's Quantitative Research Group.
Data is not weighted. Use caution when interpreting small sample sizes

Figure 4: Data silos per organization are increasing

In addition to data in cloud services, consider legacy onprem servers that host mission critical applications that cannot be retired.. An IBM mainframe with internal storage that runs chemical lab equipment, legacy Unix servers attached to small storage arrays, and virtualized server farms hooked to large storage arrays are all legacy infrastructure that is still critical to many businesses today. In addition, there is likely data in different platforms, accessible by multiple OSes—legacy Unix and Windows, Linux, macOS—and a mix of database formats—SQL, Oracle, even NonSQL. Managing, and protecting this disparate data, in addition to having a backup and recovery strategy to address it, is a serious challenge. Indeed, Rubrik's "The Data-Forward Enterprise" survey for data growth indicated that IT managers expect a 46% Compound Annual Growth Rate (CAGR), or a doubling of data every two years. **Figure 4**, from that same survey, shows 65% of respondents indicated they have at least 11 data silos and 15% have 51-100 or more. These data silos must be addressed in a backup and recovery plan.

Data Sprawl is not only challenging but it inhibits digital transformation (DX) initiatives using data management and analytics, since collecting and correlating the data is difficult with many silos. Moving to a cloud-based platform will help aggregate the data in fewer locations. In addition, a Public Cloud solution will aggregate data across geographic boundaries enabling convenient access to that data by home-based users.

It should be obvious that a move to the cloud is in most organizations' future, and it certainly has advantages. However, it also has its share of challenges. Perhaps the most important and glaring challenge is that of Backup and Recovery.

Backup and Recovery Strategy to Address Modern Environments

Electronic data is perhaps the single most important resource to any company. Financial, security, human resources, and product information as well as customer's personal information is the organization's lifeblood. It stands to reason that protecting that data must be a key goal of every organization If data becomes unavailable, for reasons ranging from power outage to hardware or network failure or ransomware, it will be expensive for the business in many ways.

Perhaps the most important, yet often overlooked, reason is the importance and criticality of backing up that data and the ability to quickly restore it in a timely fashion ensuring business continuity. Developing a sound, effective backup and recovery strategy demands a holistic approach.

Noting the data silos mentioned previously, most silos also have their own backup solution. Backup products evolve over time and in the case of a large organization, backup products may be as broad as an entire enterprise or applied to a single server or storage array. In some cases, organizational silos may have their own IT department that contain multiple backup products and strategies. A public cloud service often has their own backup product but these have limitations including making yet another silo for backup and recovery. In a modern environment a holistic approach will develop a standard company data backup and restoration strategy and a single solution that includes data in the cloud. Employing the shared responsibility model of the cloud, organizations should take control of backup and recovery operations for cloud data as well as other sources.

It is also important to remember that there is no use backing up data if it cannot be restored. It is amazing how many organizations do not regularly test their recovery strategy. Data recovery failure stories are plentiful. Loss of data, and downtime in restoring dataresult in loss of revenue and increased costs through loss of productivity, resulting in management frustration. Manual restoration processes are long and tedious, resulting in employees not being able to do their jobs and IT staff being used to do it as opposed to a modern, automated process.

RTO/RPO Metrics Determine Efficient Recovery Strategy

To determine efficient data recovery, it is common to use the Recovery Point Objective and the Recovery Time Objective (RPO and RTO). These are ITIL defined terms that are goals or targeted objectives relating to data recovery, in terms of restoring business continuity.

Recovery Time Objective (RTO) or Recovery Period

RTO is the duration of time and service level that a business process must be restored after a disaster or disruption of service. This determines when business continuity will be restored. It is the maximum recovery period allowed for a service to be down or access to a resource to be lost. For instance, if the RTO of a SQL database is 1 hour, and a server crashes, then the applications must have access within one hour.

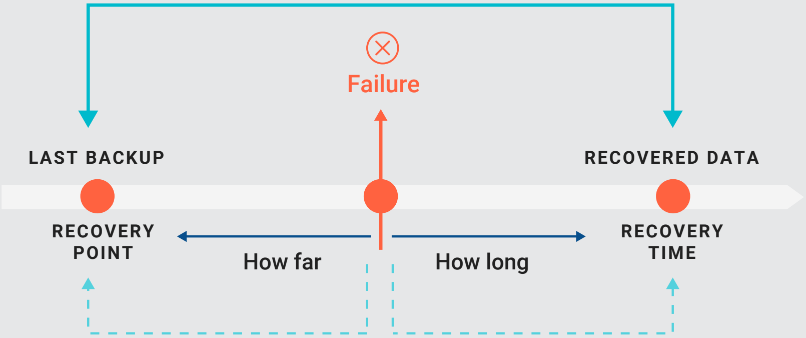

Ensuring data protection and availability requires the utilization of modern technology and services, including applicable Cloud offerings to reduce data silos. Developing a sound backup and recovery strategy to protect that data and ensure its availability to users, applications, and services will ensure acceptable business continuity goals. It is critical to use best practices to avoid data loss and productivity time by defining RTO and RPO metrics and designing a strong supportive infrastructure.

RECOVERY POINT OBJECTIVE (RPO) OR DATA LOSS

RPO is the maximum period that data access and transactions will be lost from an IT service as the result of a disaster or disruption of service. This is a measure of data loss during a disruption This is based on frequency of backups. For example, if a database has daily bakcups the RPO would be 24 hours – meaning it is acceptable to lose 24 hours of data in the event of a failure.

RPO and RTO are fundamental metrics in any backup and recovery strategy. These are limits, defined by business entities and service providers, within which business continuity can be established. It is the margin of error for the restoration of data and service. These metrics will drive the IT infrastructure's effort to support the backup and recovery strategy. This may include items such as:

- Local data backup vs offsite backup data storage
- Backup media (tape, disk, network storage)
- Network bandwidth
- Automation
- Backup power generator on site
- Disaster Recovery Plan

Service-Level Agreements (SLAs)

SLAs are not new to the IT community but are taking on new importance in modern IT infrastructures. An SLA is basically a contract between a user or customer and a service provider. There are three basic types of SLAs:

- Service level based – between one service and all customers of that serviceCustomer based – between a single customer and all services used

- Multi-Level - SLAs vary by departments, organizations and subordinate servcies of the customer.

A hardware support SLA, for example, defines how a support provider is to remedy hardware failures and get the device back online. SLAs differ depending on the business need. For instance, a server hosting a mission critical database may have an SLA of four hours return to service, while a server hosting user applications may have an SLA of 24 hours, or even several days, return to service.

EXECUTIVE CORNER

Case Study: A Million Dollars an Hour

Several years ago, a large U.S.-based company opted to purchase a new, state-of-the-art server to drive a high-end database. This server was the first of its kind the manufacturer installed at a customer site, but their customer was fine with being on the leading edge to get the technological benefits. Unfortunately, a few days after going live, the server crashed. The manufacturer's support organization dove into it, but it was not easy. At one point an engineer was writing boot code on the fly just to boot the server. The server was down over two days and the CIO was livid—famously complaining that he was losing "a million dollars an hour," reasoning that the entire business depended on access to this database and no one in the company could work without it. While he may or may not have lost a million dollars an hour, it is amazing that he did not plan for this disaster. There was no duplicate copy of the live data, there was no plan for business continuity in the event of a somewhat experimental server failing.

SLAs dictate to the support organization the resources they need to have ready, such as people, software, tools, and processes, such as monitoring and helpdesk operations. The four-hour SLA, in this case, will understandably be more expensive than the 24 hour SLA.

Note also the customer may be external or internal. The business unit may define the SLA for a database being returned to operation based on the need for the business. The business unit, then, is the customer and IT would be the provider.

SLAs should clearly state the customer, the service provider, the service provided, and observable metrics that can be measured for compliance. Some SLAs may also include penalties and may have expiration or termination requirements. Both parties must agree to the terms of the SLA

Examples of SLAs are endless but a few typical ones include:

- Hardware and Software support and service restoration – agreement in time required to return a service or device to proper operating condition for business continuity
- Data Recovery and Restoration—RTO and RPO – agreement in time to restore to operation and amount of data than can be lost
- ITSM—Ticket Management uses SLAs to define time to respond, resolve and repair a software or hardware incident.
- Cloud Services
    - Deployment of servers, such as Virtual Machines for IT and business projects
    - Deployment of software
    - Security services, such as antivirus metrics
    - Data restoration

The important point about SLAs is that they drive IT operations. The provider must do what is required to meet the SLA—meaning employing more people, incorporating new software or hardware, faster network speeds, and other services to meet the SLA's termsDefine the SLA to meet business needs, and that will drive the provider to meet those requirements.

One way to ensure that data protection SLAs are being met is with a *declarative policy engine*. These are provided by some products. Most administrators are familiar with the imperative model, which is a clearly defined step by step process listing tasks that have dependencies. Those tasks might include:

- Define the workload – machine or data source
- Data compression, deduplication
- Define backup job priority.
- Define backup time, frequency
- Connect to storage target
- Validate backed up data

The problem with the imperitive model is that it is very time consuming and it is a house of cards, each task depending on the previous one. If one task fails, it all fails. For instance if the data compression task fails, the backup will fail. If the connection to the target fails, the backup fails, requiring manual intervention.

In the declarative model, the end state is defined and the tool does the work. This is a powerful concept, allowing the SLA to truly drive the end state - such as the RPO. In addition, the declarative model can apply the SLA to all workloads—bare metal, virtual, SQL and Oracle databases, and more. A powerful declarative engine keeps the administrator out of the mundane, time consuming task of scripting or manipulating tools for various workloads.

CHAPTER 2

Secure Data and Infrastructure

In This Chapter:

- Defense in Depth Strategy
- Achieving Immutable Data
- Mitigating Ransomware and Malware Attacks

Ransomware and malware proliferation will continue. Businesses need to focus on eliminating attack vectors through lifecycle management, for example, patching, and using a "Defense in Depth" approach. Even then, ransomware or a myriad of other security issues can still occur. So, businesses need to have a strong recovery plan as well as a proactive security methodology.

DEFENSE IN DEPTH STRATEGY

Defense in Depth Strategy was developed by the National Security Agency (NSA) to be an to information systems security comprehensive solution. It requires multiple, redundant defensive operations to mitigate security failure by reducing the attack surface.

A Defense in Depth Strategy

A Defense in Depth Strategy is any solution that fights ransomware, malware, viruses, or other cyberattacks . This is an overarching security strategy that includes many components.

There are three levels of control in the Defense in Depth Strategy -physical, technical and administrative.

Level 1 Physical Controls—Examples include door locks, employee access badges, biometrics, CCTV systems, and security guards or even guard dogs. These controls limit access to physical facilities. You may be surprised at how easy it is to gain physical access by tailgating an employee into the building, employees forgetting to lock the door, and similar actions. One example includes a small business which had their server located just inside the back door. The employee who opened in the morning found the door unlocked on many occasions. It would take 5 minutes to steal this server, which had health records on it.

Level 2 Technical Controls—These controls include disk encryption, fingerprint readers, facial recognition, authentication (login), authorization (read/write access to files and folders), antivirus and ransomware protection software, and other technical limits to access systems and contents. As noted before, these can be easily compromised—hack a password, trick a user, transfer a fingerprint, etc. Ransomware protection software is not perfect, but with that said, there are advanced controls that can mitigate attacks.

Level 3 Administrative Controls—This includes implementing policies and procedures such as hiring practices (background checks), data handling procedures, security requirements and service level agreements (SLAs).

A good Defense in Depth Strategy should include:

- Fast, efficient detection, protection, notification and removal of viruses, malware, ransomware
- Employee training to mitigate social engineering attacks such as phishing, pretexting and tailgating attacks
- Employee Training for understand security policies
- Security controls specifically designed for remote employees, business partners, and supply chains
- Remove unsecure systems or networks
- Robust access security policy (passwords, resource access controls)
- Enforce security policies
- Apply current patch updates in a timely fashion

State of the Art Defense Mechanisms

Today's modern, complex IT infrastructures require equally modern defense tactics such as *data immutability* and the *Zero Trust Network.*

Achieving Immutable Data

Providing immutable protection for company data will protect it against any attacker. It is important to understand what "immutable data" is and why it is important to a backup and recovery best practice strategy.

DATA IMMUTABILITY

Data Immutability is simply the concept that data cannot be modified (read/write/change/delete) after it is created.

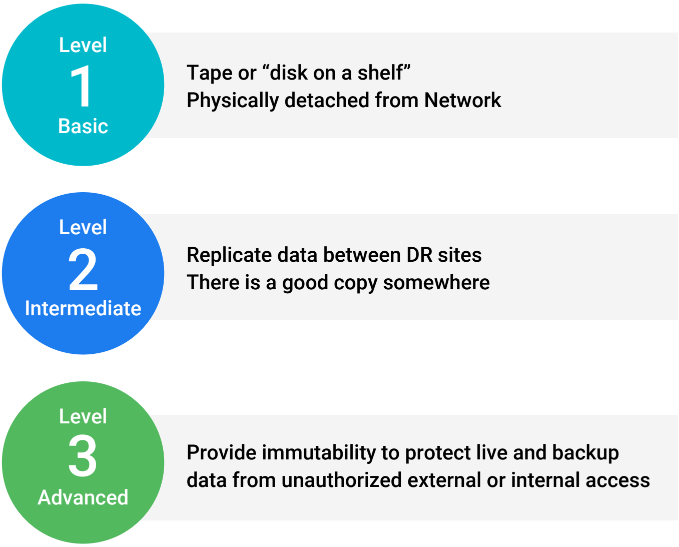

Level 1 Basic

Tape or “disk on a shelf”
Physically detached from Network

Level 2 Intermediate

Replicate data between DR sites
There is a good copy somewhere

Level 3 Advanced

Provide immutability to protect live and backup data from unauthorized external or internal access

Figure 5: The three levels of data immutability

If data were perfectly immutable, that is, it could not be modified or deleted, there would be no data breaches, no credit card information or identities stolen, and no ransomware, malware, or viruses. Somehow, we must find a way to protect data from being changed or deleted by attackers and permit it for valid users.

Over the years there have been many products and techniques to safeguard backed up data. Some are valid, but many are myths. Looking at **Figure 5**, what’s your company’s data immutability level?

Level 1 Tape or Disk on a shelf or Security by obscurity—Probably every organization does this to a degree. Many believe that tape is immutable—that is you cannot read or write to the data without mounting the tape. Store it in a mountain cave somewhere and ransomware will never find it. Likewise, many small businesses pull the disk array from the storage device and lock it in a closet. No network access, no ransomware, no problem; until a ransomware attack occurs. At that point, putting the media back in the network where ransomware is still living and then that data is attacked. Even if you purged ransomware from the network and restored from tape,, tapes deteriorate over time and restoring Terabytes of data from tape may take weeks.

Disks can be mounted and still rely on the data being read-only. When they are mounted for a backup, corrupt data can be copied onto the backups, breaking the immutability. Even worse, you may not know when the corruption entered your backups. This would require detective work before a restore can occur, resulting in breaches to your negotiated RTOs and/or RPOs.

Level 2 Replicate data between sites—This is a very common data recovery strategy today, especially with databases. All the data is live and replicated to different geographic sites. An attacker encrypts data at one site—you still have two copies. This strategy hinges on ransomware not being able to go across network links or through database connectors. While these methods reduce risk, they are not immutable because:

- Experts have seen ransomware go across network links and connectors and infect remote sites.
- Infected data at site one will eventually infect other sites, unless the attack is discovered and data is purged and replicated back. Some companies have reported having ransomware infections for months, and, as explained in level 1, breaches to negotiated RTOs and/or RPOs will also occur here.

Level 3 True Immutability—There are many myths associated with immutability tiers. Overall, the historic security provided by Microsoft's Windows file system and the NFS file system is not immutable. The ACL, or access control list, method provides Authentication (user must log in) and Authorization (provides read/write access to files, folders, services, etc.). The problem is that all a hacker has to do is get access to one account and they are in the network. All they need to do is crack an administrator password, change it, and the path is clear to pillage.

There are bolt-on solutions that provide "immutabily" by using data in native format and rely on Role Based Access Control (RBAC) and file system security like SMB or NFSv3 (both of which are easily compromised). A better solution is to provide immutability where the file system was written to provide this capability without bolt-ons. This method provides a very low security risk..

True native immutability would protect the data from changes and put a shell around it to not let anyone in—internal or external to make changes. Valid users would gain access in a secure way, such as a secure API or other method. This limits the attack footprint by eliminating an attacker from attacking data simply because they were able to login to a network account. By protecting the live data, the backup is protected. Not only does the API buffer the attack, but it can provide options for your immutable data to be versioned and organized for recovery. Details about the data can be presented and analyzed in more detail during future disaster recovery workshops.

The Zero Trust Network

A recent concept, the *Zero Trust Network* or *Zero Trust Strategy*, makes the bold assumption that all users, applications and devices, internal and external, are NOT trusted simply because they were authenticated via a login. This approach requires continuous monitoring and validating of security credentials as the user accesses network resources. It is analogous to having a GPS tracker on the user and making sure they are going where they should be going, then authenticating, authorizing, and encrypting each access.

In addition, Zero Trust uses a least privilege access model. This means that users are granted the least level of access necessary for their roles. This will minimize the attack surface. It also incorporates a "Micro segmentation" method, defining micro-perimeters which prevents unauthorized lateral movement. These methods help to limit an attacker wandering unchecked through the network. If there is no reason for this person to be in an area, they are blocked.

The Snowden Attack

The Edward Snowden attack at the NSA would likely not have occurred in a Zero Trust network. An internal, authorized user, Snowden was able to download top secret material that he had no reason to access. With the least privilege principle truly applied, his activities would have been easily discovered or prevented.

To achieve Zero Trust, the organization, Crowdstrike.com; recommends the following actions:

1. Assess the organization. Define the protect surface and assess the current security strategy and tools. Then identify any gaps and remediate. Verify all default access controls. There is no such thing as a trusted source.
2. Create a directory of all assets and map the transaction flows.
3. Incorporate multiple preventative measures.
   a. Multifactor authentication. The number of authentication factors required is directly proportional to preventing unauthorized access
   b. Least Privilege access
   c. Micro segmentation
4. Provide real-time monitoring to identify attackers quickly. This narrows the window between initial and subsequent attacks. This will typically require some level of automation.

Zero Trust provides additional security for protecting data from attacks, and should be included as part of any backup solution to further mitigate risk. Zero Trust is a component of the Defense in Depth strategy.

What Is Ransomware and How Does It Work?

Viruses and malware have afflicted computer users for decades, and there is no end in sight. Perhaps the most pervasive, and expensive, is that of ransomware. Some estimates claim a ransomware attack on a business every 11 seconds, and the costs will exceed $20 billion by 2021. In fact, 2020 was the most profitable year yet for ransomware. One ransomware alone, Ryuk, netted over $3M in 2020.

As long as there is profit in attacking business and personal systems, there will be attacks and the attacks will come upon the most profitable and most vulnerable. One brand of ransomware, NetWalker, targeted the healthcare industry and sadly took advantage of the COVID-19 pandemic in 2020. Some analysts believe that healthcare is vulnerable because the data is so valuable, and it is often unprotected. Government, law enforcement agencies and universities are also popular targets. Recently an entire school district in Georgia was held up by a ransomware, closing school for days. **Figure 6** shows ransomware attackes by industry in 2020.

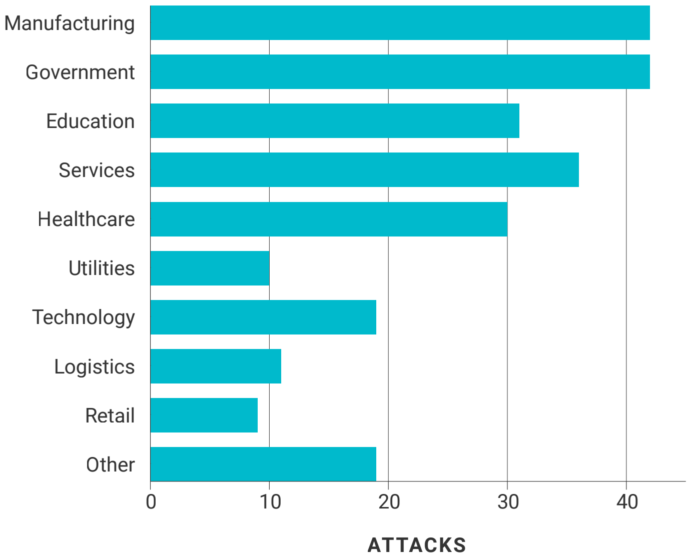

Figure 6: Ransomware attacks by industry

A form of malware, ransomware literally holds your data hostage until you pay a ransom in untraceable bitcoin. It infiltrates the computer environment like any virus and encrypts the victim's files until the ransom is paid. The most common attack vector is phishing spam through email attachments. Others exploit security holes in operating systems and other software. Once the files are encrypted, they cannot be accessed without a mathematical encryption key known only by the attacker. The user must pay a fee to the attacker.

Mitigating against Malware or Ransomware Attacks

Ransomware also attacks backups. Once it encrypts live data, the backup process will copy the encrypted data to the backup store. After a backup cycle, the live data and backup data are both encrypted. Ransomware attackers understand that restoring from backup is a good defense to avoid ransom payments and have developed the capability to go after the backup as well. Thus native immutability discussed previously, is the best defense as it natively protects the data no matter where it is.

How to Defend Against Ransomware

The first line of defense against ransomware is to design and implement a proper Defense in Depth strategy, as noted previously. Especially focus on:

- Apply operating system patches and updates in a timely fashion. Many viruses are written once a vulnerability is discovered and the patch has been announced by the vendor. The virus is written for that vulnerability, making it dangerous only to those who do not apply the patch.

- Apply best practices to guard administrator and other privileged accounts, including those required by new software, software under development, or in non-production environments. Change the default account name and password—hackers know these and test them first.

- Implement strong passwords or even Passwordless authentication.
- Implement multi-factor authentication (MFA) such as biometric authentication.
- Train users to avoid phishing and other social engineering attacks. Install antivirus software. Unfortunately, most AV software products have ransomware protection only as an add-on, and some don't offer it at all.
- Employ effective backup and recovery software and strategies. Installing backup software alone does not protect anything. Best practices will be identified throughout this guide.
- Develop immutable backups as much as possible.

One User

One former hacker turned security expert said all a hacker needs to do is trick one user in the company into downloading something. One user downloading an email attachment, or even a new game, is all it takes.

Recovering from a Ransomware or Malware Attack

It is a reasonable strategy to assume you will be attacked by a ransomware infection and plan accordingly. Do not assume current strategies are protecting you. You may just be lucky, or you may be infected and not know it yet.

The primary way to recover from a ransomware attack, and to hopefully avoid paying the ransom is to:

1. Be prepared —have immutable backups as much as possible, with a fast, proven recovery method. Define SLAs for RPO and RTO and design methods and infrastructure to support them. Incorporate modern, automated tools and methods to support the SLAs. Manually mounting disks and tapes and restoring data could take days.

2. Be vigilant. Monitor the environment and use automated tools and sound methods to keep intruders off your network. Many ransomware victims admit the attacking software was in their network for months before the attack.

3. In the event of an attack:

    a. Determine the "blast radius"—identify affected systems, files, and so forth.

    b. Identify what needs to be restored and where it should be located (local server, cloud system, or even user laptops)

    c. Restore the data, again using modern, automated tools and techniques. A valid Defense in Depth strategy will ensure achievement of defined RPOs and RTOs.

The importance of immutability in the backup cannot be overemphasized. A popular myth is perpetuated that you cannot attack that which you cannot find. You cannot depend on hiding it. If the disk or tape is locked in a vault, as soon as it comes back on line it will be attacked. If it is in a remote, locked down site, replication of encrypted data will wipe it out. However, with true immutability, even if an attacker found the data and gained access, they would be prevented from changing or encrypting that data.. If the data is truly immutable, meaning there is no write access to anyone, then no matter who the attacker imitates, they will not be able to access the data.

A new attack being spread takes advantage of reading documents – trade secrets, medical or financial data – and threatening to expose them to the public or sell them. Thus immutability solutions that permit read access are vulnerable. Storing data in a non-native, encrypted format will protect against these types of attacks.

Assuming an attacker does damage to live data and systems, the successful recovery of data to defined RPO/RTOs is dependent upon using modern, automated tools, technology, and processes. Chapter 5 will expand on the importance of automation.

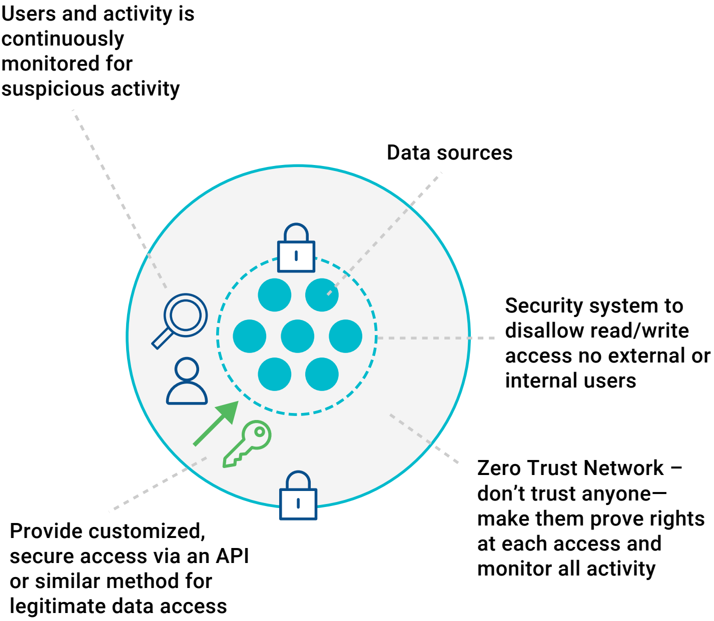
Users and activity is continuously monitored for suspicious activity
Data sources
Security system to disallow read/write access no external or internal users
Zero Trust Network – don’t trust anyone— make them prove rights at each access and monitor all activity
Provide customized, secure access via an API or similar method for legitimate data access

Figure 7: The immutable, zero trust, security model

Security from a Holistic Point of View

Let's review by looking at security in a holistic sense in order to mitigate and recover from malware, ransomware and other attacks. We have established:

- Ransomware will continue to wreak technical and financial havoc in all industries and is increasing exponentially.
- A sound Defense in Depth strategy must be defined and implemented.
- Data backups ideally should be provided with immutable data protection.
- Adopting a Zero Trust network (**Figure 7**) will deter or prevent unauthorized attackers access to the network.
- Protect business continuity by adopting modern, automated methods to restore lost or corrupted data in a manner that will achieve defined RPOs and RTOs to.

CHAPTER 3

Chapter 3 Expand Remote Work & the Hybrid Cloud

In This Chapter:

- IT challenges of a remote workforce
- Data Governance
- The Hybrid Cloud

If businesses have not already looked at the public cloud to enable a remote workforce, it is just a matter of time before they are tasked with doing so. IT organizations need to be prepared and have a strategy to conduct some, or even all, of their operations in the cloud.

IT Challenges of a Remote Workforce

The remote workforce, while it has been part of the corporate network for many years, has expanded exponentially due to the worldwide COVID-19 pandemic. The pandemic forced most businesses, large and small, to close offices and have employees work remotely. One large public utility in the U.S. will not allow employees in the building without advance executive approval, and they anticipate this environment until late 2021.

Consider the impact on business in 2020:

- According to Flexjobs.com, 4.7 Million employees (3.4% of the US workforce) were working from home before the COVID-19 pandemic.
- It is forecasted that nearly 30% of the workforce will be working remotely at least several days per week by end of 2021. **That is a 900% increase in 24 months after the pandemic hit!**
- Gartner Inc. reported that 74% of companies plan to shift some of their employees to work from home permanently, and 4% plan to move 50% of their employees to working remotely.
- 77% of remote employees say they are more productive working from home (CoSo Cloud report)
- U.S. companies who permit working from home reported having a 25% lower turnover rate.
- **54% of IT professionals claim remote workers are a greater risk.**

We can conclude from these statistics that:

1. The remote workforce has been useful for businesses and popular for employees, but the COVID-19 pandemic has caused a huge increase in people working remotely —an estimated 900% explosion from pre-COVID by end of 2021.
2. Company advantages include employee satisfaction and increased productivity, savings on office space, services and equipment costs, and lower turnover rate.
3. Company disadvantages include greater risks and challenges:
    a. Lack of control over infrastructure
    b. Greater risk to data leaks
    c. Larger virus and malware attack surface

d. More difficult to control employee behavior

e. Geographically dispersed data management and control

f. More activity from homebased workers increases complexity regarding backup and restore strategies

g. Complexity is compounded due to large geographies, with more dispersed and unreliable networks, and an increasing number of endpoints

Rate your workforce as to how the new, expanded remote workforce dynamic affects your IT processes using **Figure 8**.

In addition, consider any anticipated increase of remote workers. A move of the employee base to more remote work locations will drive decisions discussed later in this guide, such as moving to cloud services and automation.

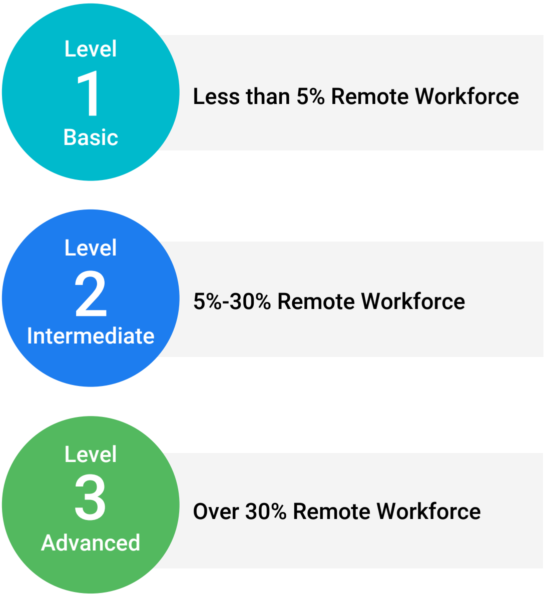

Figure 8: The three levels of the expanded remote workforce

Data Governance

With IT infrastructures, associated data locations, management, and security being put at risk with an increasing remote and mobile workforce, it is important for IT management to proactively define a *data governance strategy*.

DATA GOVERNANCE

Data Governance (DG) a component of data management that defines how data is handled in the organization. It is a set of principles and practices that ensure a high standard of data quality, including data controls, roles, policies, processes, tools, and measurements to govern the data lifecycle.

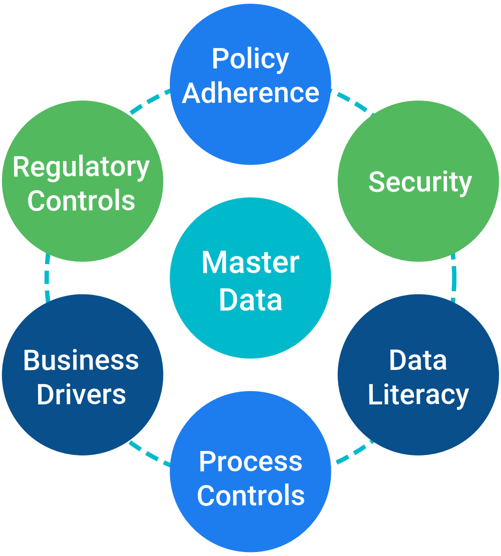

Figure 9: Data Governance principles

Normally associated with Big Data, data governance principles can be applied to any organization. These consist of several key goals, as shown in **Figure 9**, and are described as follows:

- Policy adherence and consistent decision making, following known processes, to create policies and standards.
- Improving and consistently evaluating data security.
- A data literacy program that increases the organization's data maturity.
- Process controls that adhere to defined policies and standards, including risk management.
- Define business drivers for data control.
    - These should be driven by C-Level executives to respond to external regulations.
- Meet regulatory compliance such as the European Union's General Data Protection Regulation (GDPR), and the California Consumer Privacy Act (CCPA)

Developing an effective DG strategy includes defining a charter that identifies the Vision, Mission Statement, Goals, Success Measures, Required Capabilities and a table of personnel roles and responsibilities . There are various Data Governance Templates that guide DG strategy development . One such template can be found at Smartsheet.com. Consider a few different template examples to identify one that fits your organizations style and needs. The Data Governance Institute is also an excellent resource for defining a data governance strategy.

For our purposes, a DG strategy will make sure the expanding data footprint is maintained and controlled in a strategically defined, secure, and resilient manner. It will drive implementation of tools and strategies such as backup and recovery and takes a holistic view to reduce data silos.

The Challenge of Geographically Dispersed Workforce

While a geographically dispersed (aka work from home) workforce has many benefits to the company and its people, as noted previously, it presents a number of IT challenges including security, data compliance, and data backup and restoration.

A global company must deal with servers, web hosting sites, databases, applications, and other data sources spread over a network that can have a vastly disparate configuration with potential slow link speeds. That cause latency between users, applications and data. Having a large remote workforce compounds the problem. In addition, managing IT staff rights and security groups is more complicated.

In addition, crossing government boundaries such as states and countries, enforces a wide range of inter-state and inter-country regulatory requirements. Laws such as the European Union's General Data Protection Regulation (GDPR), and the California Consumer Privacy Act (CCPA) make regulatory compliance more important. The GDPR regulates how companies protect EU citizens' personal data, and the (CCPA) is a California state law that puts more stringent requirements on businesses to protect consumer privacy rights. These and other laws regulate personal information collected by businesses and puts protecting consumer data responsibility on every business. In addition, many countries have strict laws on data crossing country borders.

Reducing and Limiting Backup/Restore Silos

Considering geographical location, IT infrastructure and privacy laws can have a profound impact on how and where data is stored and recovered. This can lead to a proliferation of backup silos with on-premises backups in multiple physical locations, private and public cloud configurations, and legacy solutions for single use servers or applications.

Developing a backup strategy is challenging when having to account for these variations and compliance requirements.

Over time, to meet the needs of business growth and security threats, an organization adds more hardware and software solutions,. This results in a proliferation of products and services in your IT infrastructure. For example, Microsoft's Windows Defender is provided free in Windows 10 and Windows Server. This is great because it protects a device out of the box, which can then be replaced with a more robust solution. However, an organization may find pockets of clients that were left with Windows Defender and not updated to the corporate virus solution.

Multiple backup solutions may also be deployed in the infrastructure. This could be caused by legacy systems having built-in solutions, and multiple cloud solutions, each with its own backup configuration. In addition, many solutions use backup agents on clients which complicates the process. Look for products that leverage APIs and native methods as noted in Chapter 5 of this book. This can include organizations, without a mandatory central IT department, having a variety of backup and antivirus solutions deployed. Multiple security and backup solutions are not only complicated to manage and administer, but it gives inconsistent results. Ransomware may be able to attack certain parts of the company where others are protected, and it may be difficult to define RPOs and RTOs for data restoration if there are multiple products to deal with.

For instance, if the solution provides storage for backing up a customer's Azure account to another account, crossing azure account boundaries, then a more complex solution needs to be designed. This is necessary in order to manage silos that are built into services, as different teams or organizations likely planned the approach for different services.

From a customer standpoint, only one solution is needed to predictably plan RTO and RPO SLAs. Ideally, a logical, straight-forward framework would be used to coordinate the recovery, security, and cost management of the solution. These are factors as egress charges and data access issues will need to be considered and managed to minimize unplanned security and financial exposure. The issue is even more complex when it crosses cloud vendors.

This is yet another case for a solid corporate Defense in Depth strategy, which would ideally limit the company to a single vendor for antivirus software and backup/recovery solutions. This would save time and money and make it easier to define business- driven metrics rather than being limited by software capabilities. That said, in large organizations there will likely always be multiple backup and security silos, especially for legacy products or due to regulatory compliance.

To build resiliency for managing multiple tools and products is to have solutions containing APIs that can be integrated in with other tools. APIs are a powerful approach which blurs the lines between products and are less impactful. Again, a good Defense in Depth strategy as well as a Data Governance Strategy will identify these silos that can present a unified solution.

The Hybrid Cloud - Getting Data Close to the User

One of the most powerful solutions for a geographically dispersed, remote workforce, as well as an efficient way to manage growth and data sprawl, is the Hybrid Cloud.

The key for enhancing productivity and performance is to get data close to the users.The Public Cloud, a concept based on commoditizing compute and allow pay as you go services, is sweeping the industry. First introduced in 2006 with Amazon Web Services, with over a dozen providers today, the public cloud allows the data to be stored in the cloud – a service that puts data closer to the users via the internet rather than connecting to on-site data centers. This is a popular way of addressing the remote workforce.

An independent IDC report indicated only 9.2% of organizations have a single, centralized data management system or platform, and that an investment in private, public, hybrid and Multicloud solutions is their highest priority. Any organization considering public cloud services should note that Gartner identified the following workloads as best suited for the cloud, and most organizations can identify some or all of these as workloads that affect their company to a high degree:

- Mobility
- Collaboration and content management
- Videoconferencing
- Virtual desktops and remote workstation management
- Scale out applications, like Data analytics and Machine Learning
- Disaster Recovery

In terms of the remote workforce, the public cloud is a powerful tool. Moving key applications to a public cloud, accompanied by defined SLAs, takes a great deal of burden off the IT organization. Just define what is needed, the SLA to achieve the need, and then write out a check! No worries about capital equipment purchase, no worries about buying extra storage hardware for expansion, and no worries about how to provide service to an expanding remote workforce. No worries about providing security, backup and restoration services and associated management. The big worry may be how it is paid for. These services are not cheap and not easy to navigate the cost structure. While not a silver bullet for all IT problems, it does have it’s place.

Obviously, there are some worries about putting services in the public cloud, and reasons to not just write a check and forget it. But for certain cases, this solution is acceptable. For example, services that can easily be put on cruise control in the cloud.

- Data widely used by remote workers.

- Data not sensitive or classified.
- Easily administered data, often needing shared administration.
- Data that is monitored centrally, but physically distributed, for example from IOT devices.
- Backup services should store data in your account. Moving data to a separate account can induce additional charges.

Keeping costs down is key in making a cloud strategy work with public and Multicloud models. A Forbes Magazine author cites a study that claims respondents estimate they are wasting 30% of their cloud expenditures and 58% indicated optimizing cloud costs is a top initiative. Cost overruns are often caused by:

- Underutilized subscriptions.
- Paying for infrastructure on discontinued projects.
- Paying premium costs when demand spikes, and not understanding and managing this carefully.
- Pricing for cloud use is typically based on complex, variable pricing methods. Combined with lack of proper tools—or interest —in tracking usage and cost overruns are expected.

Hybrid Cloud Recovery Services Challenges

Perhaps the most interesting item in Gartner's list of best Public Cloud workloads , cited previously, is Data Recovery. Even in a disparate, largely remote environment, it is easy to perform backups. Data Recovery, however, is a different, and more difficult story. Gartner predicts that 50% of organizations will increase their budgets for cloud-based disaster recovery (DR) solutions by 2023.

Building an effective data restoration strategy requires achieving desired RPO and RTOs defined by the business entities. This includes everything from restoring a few files or folders to Terabytes of lost or corrupt data. Equally as important is protecting the data from ransomware or Malware attacks. Remember that mitigating risk for ransomware attacks is centered around the ability to quickly recover data with minimal loss. To accomplish this, IT managers should provide a backup solution that will allow recovery for legacy systems as well as cloud-based data.

Table 1 shows examples of how a data restoration strategy might be accomplished. Make sure each type of system has a service and proper security to protect it.

Systems	Service	Security	Comments
Legacy	Backup to Tape/Disk	Offsite storage	Long recovery times for large data repositories
Legacy and Private Cloud	Multi-site Replication	There is a good copy somewhere Immutable data protection	Not a guaranteed recovery method vs ransomware
Public and hybrid cloud	Cloud Backup—provided by Cloud services	Dependent on cloud provider Immutable Data protection	Short RTO, RPO required
DRaaS	Disaster Recovery as a service from Cloud provider	Dependent on cloud provider—Immutable data protection	Long RTO, RPO

Table 1: Data Recovery Restoration Strategy

DR as a Service typically cannot provide immediate recovery since it usually needs to be converted to another instance and associated overhead. It is effective but not instantaneous

In the current movement to a remote workforce, with a geographically dispersed data location model, when increasing migration to Cloud services it is critical to design a Data Governance plan, analyze current and future benefits in the cloud, and design a backup and data recovery solution to fit the environment that is resilient enough to expand to future needs.

CHAPTER 4

Prepare for New Workloads and Apps

In This Chapter:

- Applications and the Cloud
- Enterprise Backup and Recovery Strategy
- Choosing a Solution

New applications are hitting the market at a breakneck pace. IT organizations need to be prepared to onboard new apps and workloads and know how to protect them, including their data. Providing resources to remote employees in a diverse geographic network requires evaluation and adoption of a hybrid cloud strategy.

In a recent report, Gartner indicated that the move from centralized to distributed work environments, brought on by the Covid-19 pandemic , has enhanced the adjustment need for data and related capabilities to support these distributed environments. Further, the report suggested that IT management must move data availability by moving assets and processing to cloud and edge environments. In other words, moving data, data assets, and applications closer to the users will improve productivity and user experiences. However, networks with low bandwidth or intermittent connectivity can seriously challenge this.

Gartner further suggested that one way to improve the remote work experience is to move data assets to the cloud . For data and services entirely contained in the Cloud, applications and related services are provided at acceptable speeds and performance levels. This solves the slow link problems, and offloads IT staff from solving these issues, making them more productive for other projects and duties. However, most organizations will have a hybrid environment with on-premises, public and private cloud data sources, making it a more complex issue. However, even a gradual move to the public cloud will gradually simplify the infrastructure.

The adoption of public cloud environments as a tool for enterprise data solutions will help meet the demands the remote workforce.. Applications using a SaaS (Software as a Service) model will be the easiest way to start moving in this direction, especially if the data can be mostly contained in the cloud. For example, if your email and inbox mostly reside in the cloud, then the amount of data actually moving to and from the cloud is minimal as you send an email or download an attachment.

Applications and the Cloud

A move to the cloud implies that things must be done differently than the legacy strategies of providing on-premises applications, storage, recovery services, and supporting infrastructure. Applications are a particular challenge due to the explosion of new applications for mobile, desktop, and even data center environments. Vendors of products and services, from consumer products to healthcare, are continually developing mobile apps to make it easy for customers to order products, get support, and provide data. A VMware blog recently stated "...it's predicted that the number of applications created in the 5 years ending 2023 will be greater than the amount built in the previous four decades."

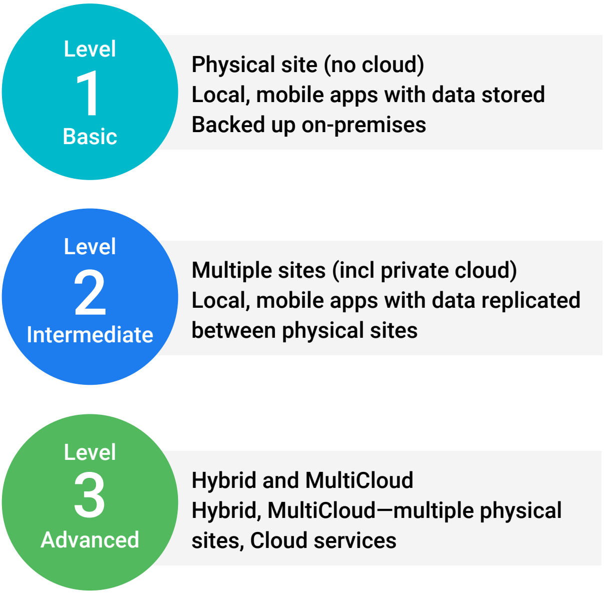

Level 1 Basic

Physical site (no cloud)
Local, mobile apps with data stored
Backed up on-premises

Level 2 Intermediate

Multiple sites (incl private cloud)
Local, mobile apps with data replicated between physical sites

Level 3 Advanced

Hybrid and MultiCloud
Hybrid, MultiCloud—multiple physical sites, Cloud services

Figure 10: The three levels of infrastructure

Accommodating Legacy Apps

Legacy applications may not be suited for the cloud, as they are ill suited for flexibility and scalability. Making a change affects multiple components.

To prepare the IT infrastructure for the new application paradigm, it is essential to survey the current environment and determine the level of the infrastructure supporting these applications (see **Figure 10**).

This analysis, together with the compelling case for moving to the cloud, will inevitably lead to consideration of if and how to move legacy apps to the cloud.

Not all applications lend themselves to the cloud. One estimate claimed that only 30-40% of large enterprise applications are in the cloud. Considerations for moving apps to the cloud include:

- Cost—external help may be required to make the app work on the cloud services platform. The value to the organization may not be worth the cost to move it, secure it, or back it up.

- Demand—Questions such as "Is the app used by a wide segment of the workforce?" or "Are they remote or local users? "should be answered.

- SLA Requirements —Can the Cloud service provide the required SLA? Perhaps your IT staff can do this more effectively.

- Can the cloud service provide necessary security? Perhaps the cost is prohibitive?

- Do you trust sensitive data to be put in the cloud? While cloud providers make it clear that they have processes that protect your data, if you follow them. However, you are also trusting that your organization can safely manage this data, using new approaches and methods.

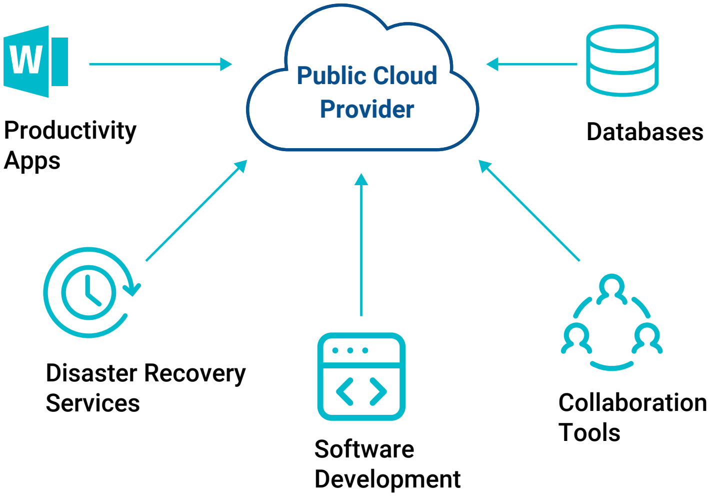

Figure 11: Applications that generally lend themselves to the cloud

Applications that generally lend themselves to the cloud are illustrated in **Figure 11** and include:

- **Software Active in Development**—cloud provides variable flex computing capacity without standing up more hardware, and it's quick to tear down.
- **Collaboration**—email, social media, and similar tools are good cloud candidates since the cloud supports work and tools from anywhere.
- **Productivity Apps**—moving to subscription services like MS Office 365 eliminates software headaches, storage, updates, access, licensing changes, and is centrally managed. It may be beneficial to invest in subscription services rather than migrating old applications.
- **Big Data and Compute Intensive Apps**—flexible storage, support and tools allow the storage and compute requirements to grow and shrink on demand.
- **Disaster Recovery (DRaaS)**—Noted previously, there are many advantages for moving DR to the cloud.DRaas providers are able to support complex environments with diverse OS and physical platforms, including VDI, and is a good solution, especially for short staffed IT department. Like any service, it is typically well orchestrated, and can handle cloud and onprem backups.

Note that there are options for moving applications to the cloud, without migrating current apps. Purchasing new subscription services, new cloud apps, and hiring migration service providers to do the work are valid options.

The Power of Cloud Apps

While identifying legacy apps to move to the cloud, don't ignore apps built for the cloud. These are not particularly new and may surprise you to learn they are "cloud apps" using a SaaS model. Some popular ones include:

- Paypal
- Slack
- Google Drive
- OneDrive
- SalesForce
- Microsoft Office 365
- Google G-Suite
- Zoom
- Zendesk

It may prove beneficial to buy into an established cloud app rather than moving a legacy application to the cloud. In fact, you are probably using some of these already. The power of cloud-based apps, whether they are commercially made, custom made, or migrated legacy apps, is the wide distribution and ease of management and use. Note, however that these apps still require configuration with data that internal business processes need.

The education world is a large user of cloud-based apps. One college IT manager described his job, at a high-tech university, as constantly trying to keep one step ahead of the students trying to hack the network, or use apps in a way that was not intended. In addition, student requirements for hardware and software change every quarter. Imagine the simplification of user support by using cloud-based apps and eliminating installing software on student laptops thousands of times a year, as well as using experts to keep the devious students out. The banner page of Brigham Young University makes the point of how easy it is from a user perspective "Work anywhere on any device and store your work in the cloud... removes the need to go to the computer lab or have a powerful computer to run software." Now every student has a link to a powerful platform without spending money on expensive personal laptops.

Cloud Application Infrastructure

It is clear that adopting a hybrid cloud structure for applicatons is the best strategy. For those that lend themselves to the public cloud, a cloud provider must be determined.

Popular cloud services include companies like Amazon Web Services (AWS), Microsoft Azure, Google Cloud Platform, Alibaba Cloud (China), IBM Cloud, and many others. Gartner again provides an excellent downloadable Cloud comparison tool that is a good resource.

There are a few basic points to consider when choosing a platform. Evaluating and comparing cost is obvious. Cloud providers vary greatly in what and how their costs are calculated. Watch for:

- Activation and termination fees.
- Pausing Workloads—ability to only pay for active workloads.
    - Some data management services
    - Intermittently use applications can be paused when not in use, but the state is saved allowing for immediate re-use without a delay for restarting.
- Compute cluster instances that are not active. Pausing saves cost of the compute resourcesData rates—by the day, by the hour—some even charge by the second.
    - Actually the smaller the increment, the better value it is.

  - For example if you pay by the hour and you only use 15 minutes, then you have wasted 45 minutes of payment.

- Penalties for overuse, peak period use, or under use.

- Volume discounts for compute heavy applications.

- Reserved capacity discounts can provide savings over pay-as-you-go.

- Scalability – This is a huge advantage for a cloud platform. Make sure you can scale storage, servers, networking and compute resources such as memory, CPU, and what platforms are available. Watch for fees associated with scalability.

Services such as Software as a Service (SaaS), Platform as a Service (PaaS), and Infrastructure as a Service (IaaS) are used in the cloud to create, host, and deploy applications. Note that these services, related software languages, and interfaces provide some flexibility to move apps between cloud providers, but some will be proprietary. When moving or creating apps in the cloud, keep in mind you may be stuck with that provider. Nevertheless, the services used in the cloud fall into these categories:

- **Software as a Service**—SaaS, often described as "on-demand software" consists of applications that are licensed and delivered on a subscription basis from a cloud service. In this model, applications are accessed via a web browser eliminating locally installed software. Office 365 is an example of SaaS. Free applications are typically deployed this way with fees charged for enhanced options.

- **Platform as a Service**—PaaS provides a development platform for companies to develop, host, and deploy apps while the cloud provider provides the infrastructure. Applications written on a given platform can be uploaded to any PaaS provider for that platform (software language). This architecture optimizes costs by efficiently using resources and providing maximum design flexibility.

However, it comes at a cost as the proprietary platforms are tied to the provider, and switching will require development.

- **Infrastructure as a Service**—IaaS services consist of APIs (Application Programming Interface) to manage the creation of virtual machines, using the appropriate hypervisor. Virtual machines are powerful entities that enable creation of host machines to enable installing software, providing backup and restore services, security, and management. Literally just a file on the disk, VMs can be started and stopped and reconfigured quickly. In addition, IaaS may provide Linux containers as an alternative to the hypervisor.

The power of applications in the cloud, no matter how you get there, will transform business demands. However, with this expansion into the cloud comes the risk of yet another data silo which must be managed, secured, and backed up. Careful planning will eliminate old silos and merge new ones. While it is probably not possible to move everything to the cloud, cloud-based infrastructure can shrink the IT footprint, reducing demands on IT staff.

Define or Update Enterprise Backup and Restoration Strategy

As applications and services move to the cloud, there is no doubt that an effective enterprise backup and restoration strategy is a successful Defense in Depth plan's chief cornerstone. This involves three key elements:

1. Defining a backup and restoration strategy

2. Implementation of that strategy

3. Testing and validation

Cloud Backup strategies

In defining a backup and restoration strategy, you must account for all elements. This includes applications, network data storage, legacy applications with local storage, private cloud integrated storage, public cloud apps, storage, and location. Performing an assessment of the environment will define the scope of what must be accomplished. Some elements might include those listed in **Table 2**.

These are not particularly best practice data points. Other information such as link speed and backup product name is helpful. In addition, it will be helpful to determine the effectiveness of those apps—are they doing the job? Is a high number of helpdesk tickets logged due to failures? Has data restoration ever been attempted and if so, what were the results? Do not wait for an attack to happen before you tighten the

Application	Classification	Deployment	Location	RPO	RTO
App1	Legacy	IBM mainframe	Chicago data center	24 hours	3 days
Office 365	Cloud	Multiple sites	AWS	Posted	
SQL Database 1	Private Cloud	Dell Cloud	Atlanta Data Center	15 minutes	8 hours
Oracle Database 2	Private Cloud	Dell Cloud	Atlanta Data Center	1hour	12 hours
SalesForce	Public Cloud	AWS - Multi region Tier	distributed	Defined – near zero	Defined – potentially zero
Banking App2	Public Cloud	AWS Pilot Light tier	Detroit sales office	Defined 30 minutes	Defined 1 hour
Department businessApp3	Legacy	HPE Server Virtual	Dallas Sales office	24 hours	24 hours

Table 2: Elements to consider when defining a backup and restoration strategy

ship. Work with business owners to define or tighten up SLAs. Those SLAs may drive the move to a new automated service.

Strategy Implementation: Choosing a Solution

Implementation of the backup and recovery strategy begins with determining if the current solution(s) truly meet the needs of the enterprise. Taking a holistic approach will identify what the current and anticipated future needs are and then find a product to meet those needs. Many organizations will find several backup solutions exist for the various data silos. Reducing those silos, as mentioned previously in this guide, will also reduce complexity of the backup/restore operation, and make it easier to find a compatible solution.

Obviously, the best solution is a single product—a one size fits all solution. However, products are selected based on "cool features", familiarity with it, or cost rather than matching features to needs. A good backup / recovery product should include:

- **Support of Legacy Systems**—Many legacy systems still use tape backup, so a solution needs to support those systems. It may also be possible to access the data for backup without using the tape system.
- **Support of Private Cloud Systems**—Private cloud systems, often a standalone or converged infrastructure, may have its own backup solution. These should be flexible enough to adapt to a new product.
- **Support of Multi-Site Datacenters**—Relying on data replication. For organizations with multiple core data center sites, the recovery strategy may be to simply use live data on-premises. "There is a good copy somewhere", but it is not impervious to ransomware attacks. The backup solution should support replication and address the possibility of ransomware encrypting all sites.
- **Support Cloud Archival For an offsite data repository**—this is ideal for organizations with a single data site. It is ideal if the

solution can recover workloads directly to the cloud, achieving a minimal RTO requirement.

- **Support of Remote Workforce**—The backup product should protect remote clients (with a client component) and geographically dispersed data locations, including applications and their data.
- **Support of Cloud Applications and Services**—For applications and data that reside in a public cloud, the cloud service will provide a total package of management services, including backup and recovery, or the client can provide their own solution. It is important to note that cloud services often backup data into a separate account repository that not only requires additional charges, but it may be sharing space with other companies' data.
- **Support of Defined SLAs**—To meet business continuity goals, the product should support the Defense in Depth strategy for defined RPO/RTO requirements, including recovering from ransomware attacks.

In preparing for or catching up with the explosion of apps and the new workloads of the present and future, it is critical to provide a backup and recovery solution to meet those needs—whether in legacy on-premises systems or advanced cloud applications.

CHAPTER 5

Automate IT Processes

In This Chapter:

- Automation as a Solution
- Adopting DevOps Processes
- Automated Disaster and Recovery Strategy

It is obvious that the modern IT environment is very complex, with legacy on-premises systems, private cloud architecture on site, public cloud services and applications, and a widely scattered remote workforce. Automation reduces risk and increases productivity. The more IT processes a business can automate, the more time they can spend on creating value in other ways. This modern environment cannot be successful and efficient without employing automation, at least to some degree.

In a complex, global, and cloud-based infrastructure, automation is key to meeting critical data recovery SLAs, in addition to other tasks like deployment and management.

While large global enterprises are more likely to adopt sophisticated automation processes, there is value to all businesses. It is interesting to note that a recent IDC report indicated two survey results that seem to conflict. While more than 60% of survey respondents rate managing Multicloud, hybrid cloud, data security and disaster recovery as “major” or “extreme” challenges, fewer than 33% of them report having fully automated means for dealing with data security, governance, backup, and disaster recovery. In other words, many know the problem, and probably understand that automation can help but lack resources, time and budget to implement automated tools and processes. However there are products that provide automation in the product.

Automation as a Solution

The explosion of applications, the remote workforce and an ever-changing workload demand has caused a spread of data that lives at the local office, in multiple data centers worldwide, and in private and public clouds. The IT organization of any company is faced with a huge challenge of backing up that data and being able to restore it fast enough to ensure minimal disruption of business operations and with minimal or no data loss. Recovering data without automation is slow and time consuming. In addition, the data must be secure to prevent attacks causing service and business disruption.

Automation, then, is key to meeting the demands of data recovery as it is software defined and takes advantage of sophisticated tools such as APIs, defining a software defined infrastructure using DevOps processes, and incorporating a sound, automated backup and recovery strategy.

The Power of the API in Automation—A Case Study

In order to demonstrate the power of automation for IT processes, let's look at the example by The Home Depot (THD), the world's largest home improvement retailer who faced many of the challenges discussed in this guide. THD has 2200 stores in North America, including Guam, and 400,000 employees.

Two key problems indicated a need to provide a more automated approach. THD had an actual RTO of 6-8 hours. Thus, key business databases could be off-line for up to 8 hours restoring data. A costly metric for a retail business. Secondly, the system registration, activation, validation, and customization of servers required 15-30 days for all 2200 sites. These problems were exacerbated by a 384K network link as well as infrastructure challenges such as:

- Siloed Architecture
- Multiple support organizations
- Multiple vendors providing CLI (command line interface) driven backup applications (read: manual labor intensive)
- Backup applications did not lend themselves to integration with larger systems
- Limited Automation being utilized other than some custom scripts

THD decided on utilizing a higher level of automation, which delivered impressive results:

- Registration and activation of all 2200 nodes in 3 days (compared to 15-30 days previously)
- RTO of 1 hour or less

The key element in automating their process was an API-centric design that allowed greater integration with existing software and services. This also enabled a single, unified solution that reduced complexity and eliminated having multiple members of the IT staff trained to do different operation tasks . This was done by leveraging the interface via the use of powerful APIs.

While THD had developed several PowerShell scripts to perform some tasks, APIs permitted higher level automation. For example, running registration tasks concurrently on all 2200 sites, finding nodes that still needed registration, obtaining location data, updating SLAs, setting replication targets, and even adding Windows and Linux hosts. These operations sequentially would take over 200 days without automating.

THD used an edge API that was remotely accessible. This API was also used by their software developer organization to interface to a platform (Cloudbolt) that deployed and managed virtual machines. This provided a way for the developers to have a "self serve" way to execute their own deployment as well as backup and restore ops in their test environment, which removed the load from the IT staff.

The Home Depot's story can be viewed online here.

Another example is that of leading telecommunications provider PCCW who used APIs for automated backup validation and verification, cloud data management, disaster recovery, anomaly detection, and enabling test deployment scenarios using production data.

The API Approach

Product vendors should make APIs a priority in their product. Every single function, click, and action taken in the UI can be accomplished by calling an underlying API endpoint. The API approach must also include extensive documentation for their customers. This approach permits integration with IT Service Management (ITSM) systems, such as Service Now, allowing automation of service requests. APIs also permit integration of separate and disparate applications, along with software platforms allowing elegant automation architectures.

Adopting DevOps Strategy

Creating automated processes, practices, and services requires a software defined infrastructure or infrastructure as code approach. Frequently referred to as DevOps, it combines software development (Dev) and IT Operations (Ops).

DevOps has become a popular term in referring to a software defined infrastructure—an environment that is automated by software with minimal human intervention. It is a way to take a holistic, consistent approach to the automated infrastructure, which previously was an assortment of isolated software programs, standards, and infrastructures existing in various fiefdoms within the organization. DevOps attempts to bring order to chaos and put a software defined infrastructure under a single umbrella of standards.

DEVOPS

Although there are several accepted definitions, Wikipedia defines DevOps' purpose as one to "shorten the systems development life cycle and provide continuous delivery with high quality software. In practicality, it is a set of practices intended to reduce time between committing a change to a system and the change being placed in normal production."

The good news is DevOps' popularity is increasing, which provides better processes, best practices, and tools. Gartner predicts that "By 2023, 60% of I&O (Infrastructure and Operations) leaders will invest in application development capabilities to support digital business innovation," and that "By 2023, 70% of organizations will deliver a shared, self-service platform for product teams, improving their application deployment frequency by 25%." Thus, the move to an automated environment is accelerating and for good reason. The Home Depot experience demonstrates how powerful automating a couple of processes is, but to extend it to the whole organization takes defined processes and planning, which DevOps provides.

The DevOps diagram in **Figure 12** shows a continuous process strategy for development, including:

- Plan
- Code

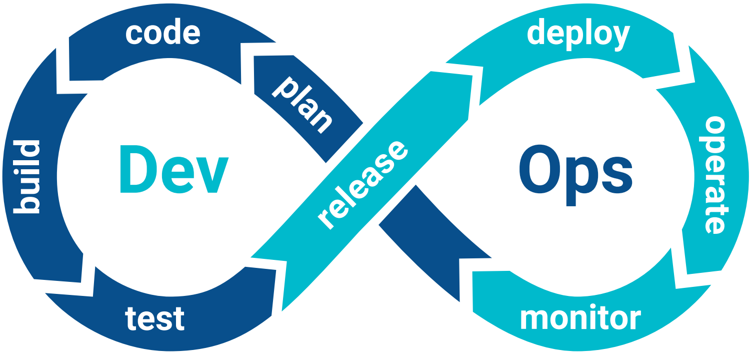

Figure 12: The DevOps principles

- Build
- Test
- Release
- Deploy
- Operate
- Monitor

The graphic here shows how DevOps is a continuous loop—not just a list of processes—it continually evolves.

Scripting Methods

DevOps is not just a set of standards—it has software configuration management tools that enable consistent automation for management, monitoring, deployment of hardware and software, ITSM, Backup and Recovery, and more. The most popular tools are *Chef, Puppet, Ansible and Terraform*.

- **Chef**—Chef's sweet spot is for deploying and managing the cloud server, storage, and software, and it uses native Ruby language - which is fairly easy to learn by a skilled developer. It uses a series of so-called cookbooks and recipes. Think of a recipe as a collection of resources—services, users, groups, files, directories, and templates. The developer can create recipes and collect them into cookbooks that provide a scripted process. Chef is a sophisticated tool, according to some analysts, best suited for AWS cloud environments.

- **Puppet**—Puppet is an open-source software product whose key feature is functioning as a vehicle for delivering, releasing, and operating software. It can define infrastructure as code, manage multiple servers and enforce system configuration.

- **Ansible**—an open source automation tool from Red Hat that is gaining popularity. It features cloud provisioning, configuration management, application deployment and other automation tasks. It is agent less and uses SSH, an thus has no custom security infrastructure. Ansible connects to compute nodes and pushes "Ansible Modules" - small programs containing the desired state to be accomplished which is then executed and then removed.

- **Terraform**—produced by Hashicorp, Terraform's sweet spot is building and changing infrastructure, or Infrastructure as Code. Terraform builds a resource graph of the infrastructure resources, then uses a planning step where it generates an execution plan. Reviewing this plan shows the end result. It enables automated changes to the infrastructure. It is helpful for building infrastructures for app configureations, creating disposable environments for testing, and spreading an infrastructure across multiple clouds for fault tolerance.

In addition, there are other tools including SaltStack and CloudFormation.. All of these are all open source except CloutFormation, which is AWS only. In terms of choosing one of these, one contributor from the Gruntwork blog offered the following considerations:

- Cloud Management vs Provisioning—examine what you will need and which tool will provide that. Some are more adept at provisioning, others at management.

- Mutable vs Immutable Infrastructure—Mutable means software updates, for example will modify the existing installation. Immutable will deploy new updated images to replace the old in their entirety.

- Procedural vs Declarative—The procedural style of coding specifies processes and actions in a step-by-step procedure, where declarative coding defines the end state and the tool determined how to get there.

- Centralilzed vs Decentralized—some tools require existence of a central server for storing state information. Communication is accomplished via a client to the central server which then pushes the updates out. A decentralized model has agents installed on each server, which periodically runs to implement updates.

- Agent vs Agentless—some tools install agents on each server, which runs as a background task for installing updates. Agentless tools typically do use agents, but they are deployed and managed by the tool itself—there is no manual intervention.

- The Community—like any other software product, these tools have user communities that can be very helpful by sharing experiences, answering questions, and even helping find hired help for a project. A Chef community, for example, shares cookbooks that can be very helpful. Consider not only the size but the strength, or activity in the community.

Consider the needs, requirements, and goals of the automation project and make sure to think ahead. Look at what the automation requirements will be in the future and find a tool that will meet those needs. It will be helpful to engage experts in this decision. For example, anyone can join the Chef community and engage in events, meetings, training and interact with the community.

In addition, custom scripting and APIs are used extensively. The Home Depot's experience is a good example. Existing scripts will likely have to be customized but it is not starting from scratch.

Choosing A Backup Solution

This clearly involves more than previously considered. The use of APIs is key to a successful automation strategy, whether it is for backup and recovery, software development, deployment or any other IT process. Thus, when choosing a vendor, consider it more of a partner—not just someone who has a cool product that checks the boxes. Besides the product features, look for a backup and recovery solution that includes:

- Standard and Open APIs. In addition to a standard vendor supplied library, some vendors offer an API repository. As discussed, customers like The Home Depot and PCCW who have developed APIs for their use could contribute them to the repository, saving time and money for others.

The Home Depot

The Home Depot automated their processes via APIs and dropped the RPO from 6-8 hours to less than an hour! It would be interesting to see the ROI on the cost of automating to get those results. Indeed, using an ROI calculation to justify the cost of products and labor to accomplish such results would be a best practice in and of itself. The company mentioned in Chapter 1 that was losing $1M per hour for down time could certainly justify costs to move to automation.

- Software Development Kit (SDK) —Most vendors do supply an SDK to assist in development of custom solutions with their product. Don't overlook this feature.

- Vendor support in API development is key in working with experts who know the product and have likely solved problems like yours already. Costs for this may vary.

The three tiers of Backup/Recovery automation are shown in **Figure 13**.

Adopting an automation strategy is key to success in protecting the modern IT environment. The case for automation, is clear, so let's see how it is accomplished.

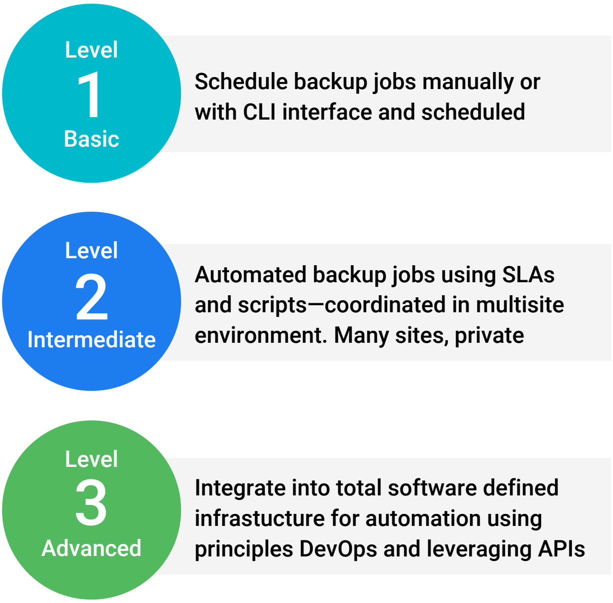

Schedule backup jobs manually or with CLI interface and scheduled

Automated backup jobs using SLAs and scripts—coordinated in multisite environment. Many sites, private

Integrate into total software defined infrastucture for automation using principles DevOps and leveraging APIs

Figure 13: The three tiers of backup/recovery automation

Automating Data Recovery

In the modern IT infrastructure, with "data living anywhere", automating data recovery is critical to an organization meeting business requirements. Overcoming physical network problems data silos, multiple backup products, and defending against ransomware attacks is difficult and nearly impossible using manual, labor intensive methods. Again, the example of The Home Depot reducing RTO from 8 hours to less than one hour is a powerful example of how automation can make a huge impact on data recovery, and an acceptable return to operation for the business even in a widely dispersed global network.

Using open APIs, in addition to other recovery automation tools and processes, provide a powerful system for reducing RTO and RPO and mitigates ransomware risk. Providing immutable live data, and a Zero Trust network with a powerful, automated backup and recovery system will help ensure business continuity.

The Time Is Now to Start the Journey

Taking the time to read and study this Gorilla Guide demonstrates your commitment to improving your company's IT processes and data recovery strategies. Identifying where you are and where you want to be is a powerful beginning, but don't stop there—it is imperative to move forward.

Study the things you've learned in these pages, and then carefully develop a plan to implement your solution. Also consider the benefits of bringing on a partner to help with your modernization and transformation efforts, allowing you to speed up the process and avoid some of the mistakes that happen in roll-your-own scenarios.

Good luck and stay safe!

ABOUT RUBRIK

Rubrik helps enterprises achieve data control to drive business resiliency, cloud mobility, and regulatory compliance. Rubrik bridges the gap between owned, on-premises infrastructure and the cloud by decoupling data from the data center through a software-defined fabric and offering a single management plane for all data, whether on-prem or in the cloud. Comprehensive data management is delivered through instant access, automated orchestration, and enterprise-class data protection and resiliency. rubrik.com

ABOUT ACTUALTECH MEDIA

ActualTech Media is a B2B tech marketing company that connects enterprise IT vendors with IT buyers through innovative lead generation programs and compelling custom content services.

ActualTech Media's team speaks to the enterprise IT audience because we've been the enterprise IT audience.

Our leadership team is stacked with former CIOs, IT managers, architects, subject matter experts and marketing professionals that help our clients spend less time explaining what their technology does and more time creating strategies that drive results.

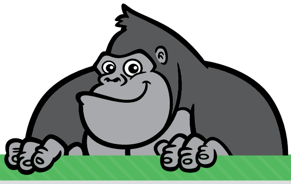

If you're an IT marketer and you'd like your own custom Gorilla Guide® title for your company, please visit https://www.gorilla.guide/custom-solutions/