Disaster Recovery to Cloud Roadmap

A COMPREHENSIVE ROADMAP FOR CALIFORNIA JUDICIAL BRANCH ENTITIES

JUDICIAL COUNCIL OF CALIFORNIA

INFORMATION TECHNOLOGY ADVISORY COMMITTEE

Contents

1. Overview 4
   1.1. Introduction 4
   1.2. Background 4
   1.3. Audience 4
   1.4. Components 4
   1.5. Project Approach 5
2. Modern Disaster Recovery Solutions 6
   2.1. Benefits 6
   2.2. Identifying Key Requirements 6
3. Cloud Disaster Recovery Components 8
   3.1. Categories 8
      3.1.1. Local Facilities 8
      3.1.2. Connectivity 8
      3.1.3. Virtual Infrastructure 10
      3.1.4. Storage 11
      3.1.5. Replication Solution 13
      3.1.6. Templates and Automation 15
   3.2. Cost Considerations 17
   3.3. Current Feasibility of Cloud Disaster Recovery 17
4. Project Methodology 19
   4.1. Phase I: Discovery & Assessment 19
      4.1.1. Key Activities 19
      4.1.2. Additional Discovery & Assessment Considerations 20
   4.2. Phase II: Selection 21
      4.2.1. Key Activities 21
      4.2.2. Additional Selection Considerations 22
   4.3. Phase III: Design 22
      4.3.1. Key Activities 22
      4.3.2. Key Activities 23
   4.4. Phase IV: Implementation & Pilot 24
      4.4.1. Key Activities 24
      4.4.2. Additional Implementation & Pilot Considerations 26
5. Implementation Examples 28
   5.1. Superior Court of California, County of Monterey (ASR / Hyper-V) 28
      5.1.1. Local Site 28
      5.1.2. Connectivity 28

5.1.3. Connectivity Costs 29
5.1.4. Cloud Infrastructure 29
5.1.5. Storage 30
5.1.6. Cloud Storage and Hosting Costs 31
5.1.7. Replication Solution 32
5.1.8. Automation 32
5.2. Superior Court of California, County of Santa Barbara (ASR / VMWare) 33
5.2.1. Local Site 33
5.2.2. Connectivity 33
5.2.3. Connectivity Costs 33
5.2.4. Virtual Infrastructure 33
5.2.5. Storage 34
5.2.6. Cloud Storage and Hosting Costs 34
5.2.7. Replication Solution 34
5.2.8. Automation 34
5.3. Superior Court of California, County of Orange (Zerto / VMWare) 35
5.3.1. Local Site (Irvine, CA) 35
5.3.2. Azure Primary Site (US GOV Arizona) 35
5.3.3. Connectivity 35
5.3.4. New Connectivity (in progress) 35
5.3.5. Connectivity Costs 35
5.3.6. Azure Infrastructure (US GOV Arizona and US GOV Texas) 36
5.3.7. Storage 37
5.3.8. Cloud Storage and Hosting Costs 38
5.3.9. Replication Solution 39
5.3.10. Automation 40
5.4. Superior Court of California, County of Placer (Rubrik / VMWare) 41
5.4.1. Local Site 41
5.4.2. Connectivity 41
5.4.3. Connectivity Costs 41
5.4.4. Storage 41
5.4.5. Replication Solution 42
5.4.6. Automation 42
Appendix A 43
Critical Court Services 43

1. Overview

1.1. Introduction

This *Disaster Recovery to Cloud (DR2C) Roadmap* is intended to serve as a reference guide for courts that are either establishing a new disaster recovery (DR) program or refreshing an existing program using emerging technologies, including a cloud-based solution. There are several considerations that impact the selection and implementation of disaster recovery solutions, such as business and operational requirements, funding, existing technology solutions, and geographic locations. Given these factors, it is neither feasible, nor necessary, for a single solution to meet the disaster recovery needs for all California courts. Therefore, this roadmap does not focus on any specific technology product or platform, but rather provides a reference for the benefits, selection, design, tools, and project methodology to support the implementation of cloud-based disaster recovery solutions.

This reference guide is intended to help courts of all sizes evaluate how modern cloud-based disaster recovery solutions align with their continuity of operations goals while also considering their current skills, resources, and investments in existing vendor eco-systems.

1.2. Background

The California Judicial Branch *Disaster Recovery Framework: A Recommendations & Framework Guide for the Judicial Branch* (available in the Information Technology Advisory Committee Library) was published by the Judicial Council in 2017 to assist courts in planning and implementing their disaster recovery strategies. To leverage this foundational work and continue to enhance the courts' disaster recovery efforts, the Information Technology Advisory Committee (ITAC) launched the Disaster Recovery Phase II Workstream. This new workstream combined efforts with a Court Innovations Program grant awarded to the Superior Court of Monterey County to pilot and demonstrate the viability of emerging cloud-based solutions for timely recovery of critical court services, and to serve as a model for interested California courts to adopt. This DR2C Roadmap is the result of those efforts.

1.3. Audience

The primary audience for this document is the court staff responsible for the planning, execution, support, and implementation of disaster recovery technology solutions. Due to the technical nature of concepts and topics related to disaster recovery systems, it is beneficial for readers to have a technical understanding of court applications and systems, networking, security, and computing infrastructure. For additional foundational disaster recovery information, please reference the *Disaster Recovery Framework*.

1.4. Components

The goal of this roadmap is to provide interested courts with information to support their pursuit of a modern disaster recovery solution, including:

- Benefits of modern recovery solutions
- Identification of key requirements to assist in evaluation and selection of cloud-based recovery solutions
- Templates and tools to deploy functional hybrid architectures that extend local priority system infrastructure to the cloud
- The viability of emerging cloud disaster recovery solutions
- Recommended project phases, tasks, and considerations for a disaster recovery project

1.5. Project Approach

When starting an initiative to implement an updated disaster recovery solution with the goal to leverage cloud services, there are four recommended project phases (see figure 1) that support a sound project methodology.

1. Discovery & Assessment
2. Selection
3. Design
4. Implementation & Pilot

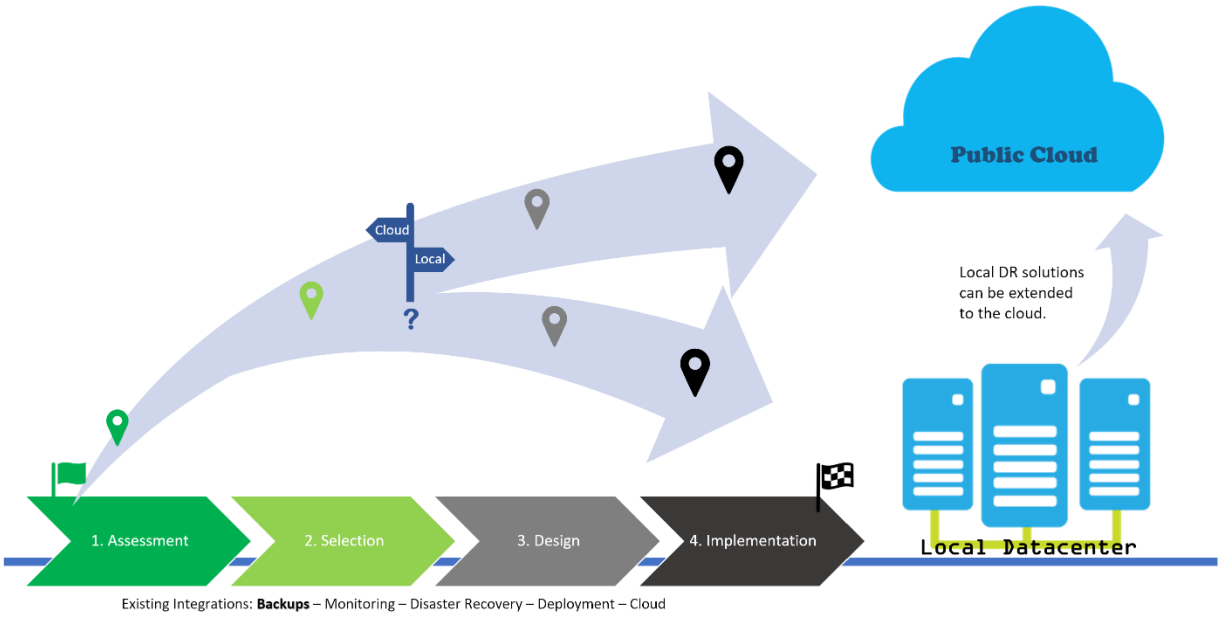

Figure 1: Disaster Recovery Project Phases

These phases of a disaster recovery project will have logical milestones along the path at which point realizing the tangible benefits of completing one or more of those milestones is possible without completing the journey in its entirety. For example, the Discovery & Assessment phase will yield detailed information about a court's local environment, which can be useful in planning for any future DR-related activities, not just in the cloud. The Selection phase should identify a desired disaster recovery solution that leverages the investments in currently used technologies to preserve local recovery capabilities while also implementing cloud replication tools.

This guide illustrates how to implement a cloud-based DR solution and the critical factors for success. It is important to note that these solutions are highly complex and require technical expertise and additional considerations to seamlessly integrate and manage hybrid local and cloud environments.

The milestones within each project phase will provide valuable information and experiences to advance a current disaster implementation, while posturing towards a future cloud migration. Detailed information on each of these phases can be found in section 4, Project Methodology of this document.

2. Modern Disaster Recovery Solutions

2.1. Benefits

With the current approach of having physical, on-premises disaster recovery solutions, organizations assume the responsibility of service disruption and recovery. In the event of a natural disaster, this dependency on physical solutions would likely result in extended down times as organizations recover from an outage. Modern recovery solutions can provide reliable, cost-effective, and powerful cloud-based computing resources to lessen the constraints of those physical dependencies and potentially shorten recovery time. The potential benefits of extending DR solutions to include the public cloud are many:

1. Reduced footprint–Public cloud offers an opportunity to reduce local resources, such as the supporting server hardware, the datacenter in which it resides, or simply the number of resources dedicated to virtualization and replicated data. Additionally, some capital investments may be reduced when moving to a cloud DR solution, which could translate to reduced maintenance costs long term.

2. Instant provisioning–Cloud infrastructure can be acquired on demand. Local DR has a hardware component that requires right-sizing, licensing, maintenance and asset replacement. Provisioning infrastructure on the cloud can be accomplished in a very short time. Cloud DR can also serve as a starting point to build out a secondary home for services as needed for those confined to a single physical data center.

3. Efficiency–Cloud infrastructure that supports the disaster recovery solution as well as the court applications that are replicated to the cloud consume only the resources needed.

4. Multi-regional protection–Cloud datacenters can be found throughout the globe and offers the ability to retain service availability in a regional or even coastal disaster scenario.

5. Service Level Agreements–Cloud disruptions are rare and can guarantee a higher level of service availability.

2.2. Identifying Key Requirements

Prior to evaluating or selecting a cloud-based disaster recovery solution, the following steps should be taken to validate a court's recovery objectives and existing systems, identify new opportunities, and cost considerations, and support capabilities for a new solution.

1. Define or confirm disaster recovery goals for *operational* continuity. These goals should be in alignment with the organizational Continuity of Operations Plan (COOP) and endorsed by leadership and operational stakeholders. This is an essential first step on which all other decisions in the process will be based. Courts looking for a good starting point to create these objectives should reference the *Disaster Recovery Framework* at: https://www.courts.ca.gov/documents/itac-dr-framework.pdf

2. Identify current backup and disaster recovery capabilities. Take time to list critical services, validate what backup and disaster recovery solutions are currently in place, and identify the currently accepted Recovery Point Objective (RPO) and Recovery Time Objective(RTO) for the organization.

3. Evaluate opportunities to simplify and migrate services and data to native cloud solutions. Determine if Software as a Service (SaaS) platforms exist that can assist with moving local services to the cloud. A common example would be the Microsoft Office 365 suite for email, document storage and personal file storage. Evaluate whether other business applications have a SaaS offering

that might make sense for the organization to pursue. The current or future plans for cloud migration are an important factor in scaling a DR solution.

4. Assess and document existing IT infrastructure and systems. Many cloud providers have discovery tools available to assist with providing a clear picture of existing local infrastructure design and dependencies that will inform the necessary cloud infrastructure design to fit a cloud DR solution.

5. Pursue solutions that leverage the current technology eco-system. It is crucial to verify that a new solution is compatible with the local infrastructure and backup/disaster recovery software already in place, as well as the local skills and capabilities for support. The solution should support both local and cloud-based replication to provide a roadmap for future expansion if a full cloud strategy cannot be pursued initially.

6. Analyze acquisition and on-going cost and budget considerations. Connectivity, infrastructure (networking, security), hardware (systems and storage), and software licensing should be included when evaluating a cloud solution.

7. Implement using a phased approach. It is easy to get excited about the possibilities that cloud solutions have to offer, however some real limitations exist. The recommendation is to take a "build-pilot-test" approach to prove the desired cloud-based recovery solution. Building the foundational components that support a hybrid infrastructure for on-premises and cloud, will allow courts to adapt to solutions that are relevant now with an opportunity to expand in the future.

3. Cloud Disaster Recovery Components

3.1. Categories

The components to consider for a cloud disaster recovery solution can be broken down into six major categories:

1. Local facilities
2. Connectivity
3. Virtual infrastructure
4. Storage
5. Replication
6. Templates and automation

This section will cover the considerations for each of these categories in more detail.

3.1.1. Local Facilities

One of the more overlooked areas of cloud disaster recovery planning is the local facilities. Even though the end goal resides in the cloud, many things need to happen to prepare the local site for cloud disaster recovery, including assessment of the current environment, existing technologies, and skills and knowledge of on-site staff. Ensuring that all these bases are covered will provide a solid foundation for the journey into successful cloud DR implementation.

3.1.2. Connectivity

The primary connectivity questions to explore are how the onsite data will make its way to and from a cloud provider and how fast the connection needs to be. Many factors can influence these answers, including:

- The desired RTO and RPO times
- The amount of data that needs to be transferred
- Whether funding is an issue
- The security requirements for the connection

Connectivity Options

Multiple ways exist to link a local side to a cloud provider to perform cloud disaster recovery (see figure 2). Although there likely are others, there are three main types of connections to consider:

1. VPN over public internet
2. Dedicated cloud provider link
3. Third-party data center link

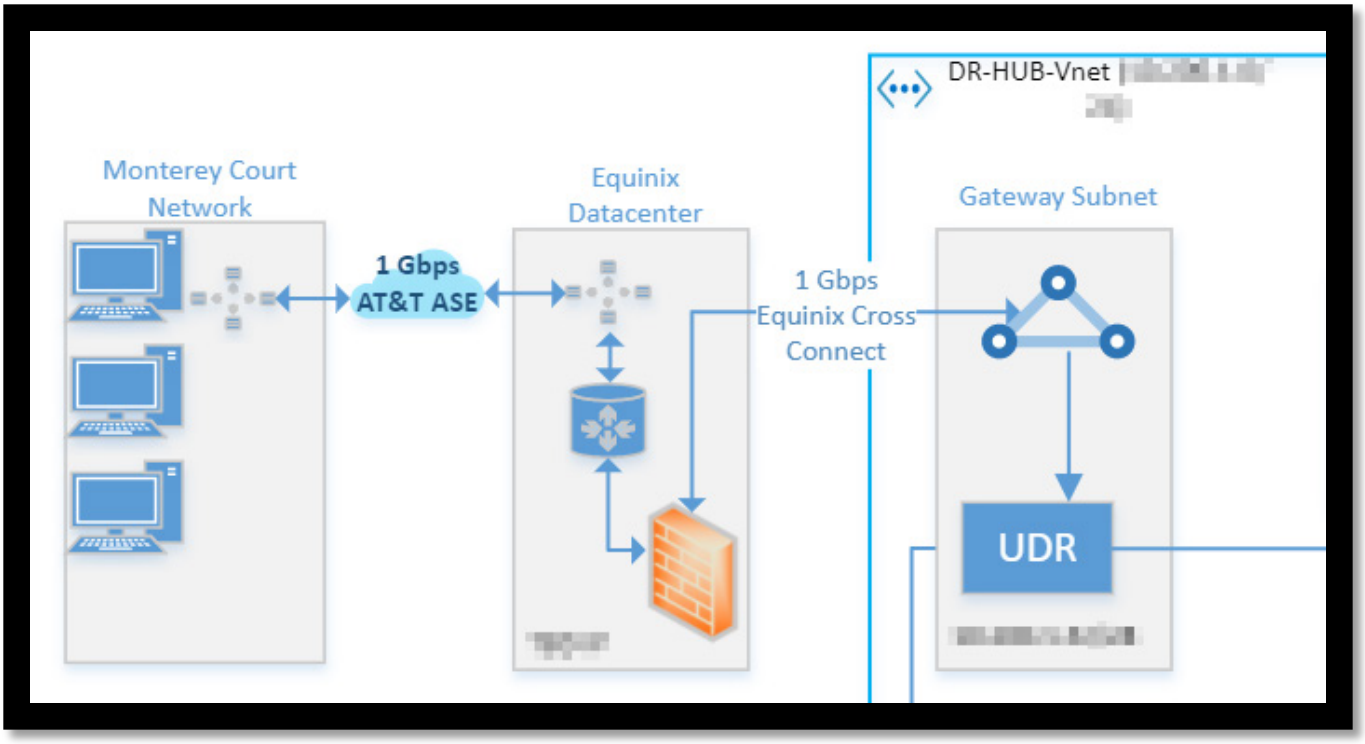

Figure 2: Sample Cloud Connectivity

1. VPN over public internet

Pros: VPN over public internet is the least complex, and in most cases, the cheapest solution to connect local infrastructure to a cloud provider. A local VPN concentrator, an internet connection, and knowledge to configure the cloud provider side are the core requirements.

Cons: Because VPNs are creating a private tunnel across the public internet, the quality of the connection is not assured and can be affected by the amount of internet bandwidth used by the local site for other activities, as well as any performance issues that may be encountered while traversing the public internet.

Note: Most cloud providers charge a monthly connection fee for VPN connectivity based on the speed.

2. Dedicated cloud provider link

Most major cloud providers such as Microsoft, Amazon, and Google now offer a dedicated private link to their cloud hosting by partnering with major network providers such as AT&T, Comcast, or Verizon. The network providers then set up a dedicated private link at an agreed upon speed to those cloud hosts for a set monthly fee. When using a dedicated cloud provider link, having an on-site resource with network expertise to configure and maintain these connections is extremely helpful.

Pros: The greatest benefit of this arrangement is a private link with guaranteed bandwidth to those cloud providers.

Cons: This solution is often very costly compared to other connectivity options. It requires technical configuration changes to local and cloud network configurations to run, adding a level of complexity.

Note: The cloud provider will also traditionally charge a monthly fee to maintain connection on its end.

3. Third-party data center link

A third option depends on geographic locations, the ability to leverage a third-party vendor's data center, infrastructure and existing connectivity for the selected cloud provider. An example of this setup includes co-location and cross-connects through Equinix to connect to Azure ExpressRoute. Once a dedicated link to the third-party datacenter is set up with network hardware hosted locally, it is possible to request a cross-connect from that hosted hardware to the cloud provider's network.

Pros: This solution can offer a substantial reduction in costs compared to the dedicated cloud provider link, without compromising any benefits.

Cons: This option is usually costlier than a standard VPN cloud connection, has more complexity than either a dedicated cloud provider link or a VPN, and requires additional hardware to be installed at the vendor's co-location facilities.

3.1.3. Virtual Infrastructure

The virtual infrastructure is the skeleton configured in a cloud provider on which all the servers of a disaster recovery solution will run (see figure 3). It also connects DR cloud storage to a cloud connection for replication, failover and failback. One of the most efficient ways to connect the cloud and the local environment is to extend the local network topology to the cloud. Because doing so connects all cloud services to each other, as well as back to local infrastructure, it is necessary to keep network security in mind when designing cloud infrastructure.

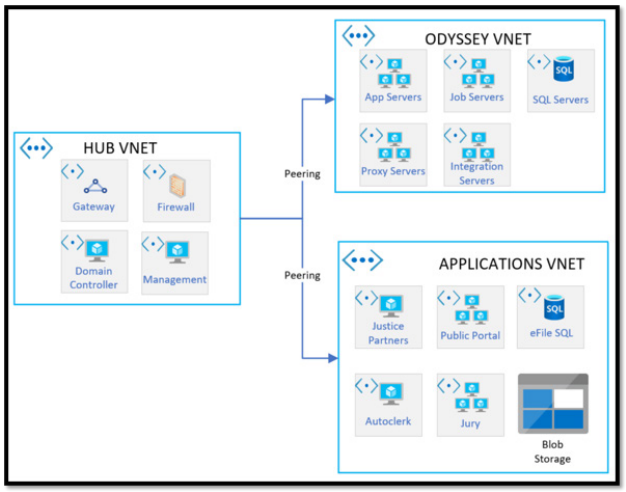

Figure 3: Sample Virtual Infrastructure

Virtual Infrastructure Considerations

When designing a cloud infrastructure, the design details rely heavily on the chosen cloud hosting provider. In this section, the virtual infrastructure relates to Microsoft Azure. However, most of the larger concepts will apply to any cloud provider. Decisions such as choosing a traditional versus hub-spoke layout or going with a hybrid cloud configuration versus keeping the cloud separate from the local site are important decisions that will also shape the design of a cloud DR solution.

1. Hub-Spoke Technology

In a hub-spoke design, the hub VNet is configured to act as the shared service connectivity housing the common infrastructure services that will manage all other environments in terms of security, access, logging, and monitoring. The hub is the intermediary point between the cloud resources and the on-premises datacenter(s).

For disaster recovery considerations, the solution will need to be self-standing. It may include but is not limited to the following services:

- Active directory services
- DHCP or IPM
- Internal and external DNS
- Internet access
- VPN access
- Perimeter security for services exposed to the public internet

These services are good candidates for the hub portion of a cloud deployment. Other VNets containing Priority 1 (P1) services act as the spoke and are connected via VNet peering.

2. Hybrid Cloud

A hybrid cloud configuration combines one or more private clouds (local sites in this case) with a public cloud provider. The benefit is a more seamless transition when moving services to and from the cloud. It requires less automation during failover/failback to configure active directory or DHCP, DNS, etc. One thing to keep in mind when considering a hybrid cloud approach is that hybrid cloud connectivity relies heavily on network design and benefits from in-house expertise.

3.1.4. Storage

Cloud storage is where all replicated data for a cloud disaster recovery solution lives. As such, it must be sized appropriately to allow data from all replicated services to be stored. Depending on the storage and replication solutions, and cloud infrastructure design, it can be one or many storage resources operating in the cloud. Storage pricing will vary depending on the availability level (cold, warm, hot; described below) and the amount of throughput up and down to the storage resource (see figure 4).

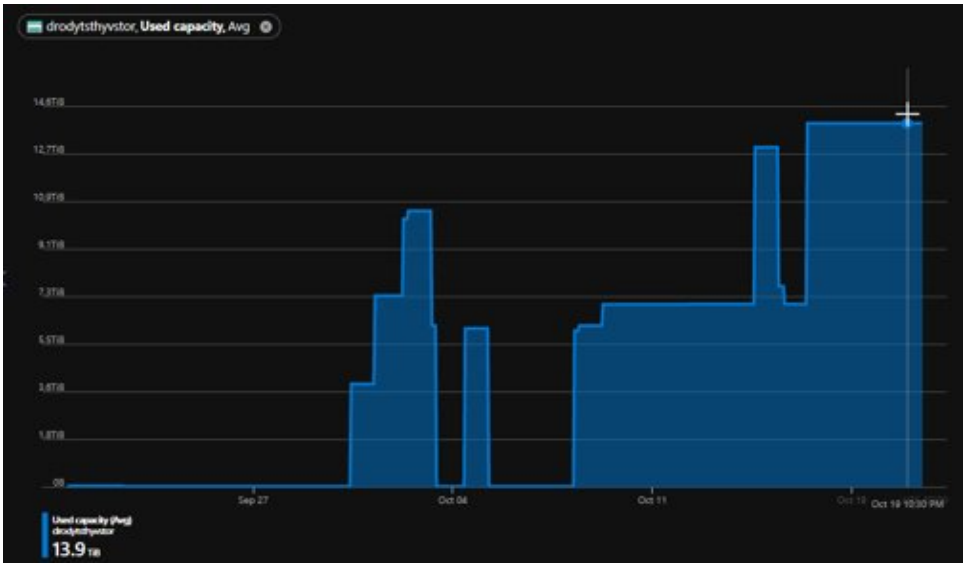

Figure 4: Sample Storage Usage

Storage Considerations

Each cloud provider has several options for replication and virtual machine data storage. Although the choices can at first appear overwhelming, factors such as the selected replication solution's preferred method, functions performed by the protected resources and the branch's financial considerations will narrow those options down considerably.

1. Hot, Warm, and Cold Storage

A given internet search produces a variety of definitions for hot, warm and cold storage. For cloud DR implementation, the most commonly used definitions are as follows:

- Hot storage is optimized for storing data that is accessed frequently. The hot access tier of storage is typically the cheapest from which to frequently access data; however, the price of hot storage itself is higher than for the warm or cold options. A live virtual machine is a good example of something that requires hot storage.

- Warm storage (sometimes referred to as "cool storage") is right at the happy medium of access costs versus price of storage. Warm storage is usually optimized for data that is accessed infrequently—maybe once a month.

- Cold storage traditionally has the lowest storage costs; however, it has the largest financial penalty to access and can take several hours to retrieve. This storage is best for archival purposes.

2. Redundancy

Many cloud providers offer different options and types of cloud redundancy. Most large cloud providers such as Google, Azure, or Amazon Web Services (AWS) operate multiple data centers worldwide. Cloud providers make locally redundant, zone-redundant, and geo-redundant storage possible:

- Locally redundant storage replicates the data within the same physical data center used to host the original data source. It is traditionally the cheapest option; however, should that physical location have an outage, data could be lost or inaccessible.

- Zone-redundant storage replicates data to multiple locations within a primary region. A region is a set of data centers that all exist within latency-defined areas and are interconnected via a dedicated high-speed connection. An example of a region would be the Western United States or the Central United States. This storage is more expensive than locally redundant storage. However, it is also more resilient than a single data center. An outage is still possible in a region but less likely than in one location.

- Geo-redundant storage replicates data to a secondary region that is typically hundreds of miles away from the primary area. It is the most durable option for cloud storage. However, it can also be the costliest.

3. Input/Output Operations per Second (IOPS)

IOPS is the benchmark of local and cloud storage options. Many cloud providers now rank their storage tiers using things like max throughput and IOPS. More expensive tiers usually have a higher max IOPS number, which will facilitate more high-performance

resources such as databases that require many transactions per second. By contrast, tiers with lower IOPS numbers are sufficient to store backups or infrequently accessed services. Backups and replication work fine between 2,000 and 6,000 IOPS, whereas something more intensive might need upwards of 20,000.

3.1.5. Replication Solution

The replication solution is the brains of the disaster recovery operation. It traditionally manages which virtual or physical machines are replicated to the cloud, the intervals for replication (affecting RPO), and handles any failover or failback processes. Considerations for this aspect of the cloud DR solution are pricing, recovery time objectives, compatibility with existing technology, and the ability to fail up and back in the acceptable timeframe.

Replication Solution Considerations

Several extremely powerful replication solutions are available on the market today. Many have similar features; however, they use different methods. In general, DR solutions can be grouped into five categories:

1. Traditional
2. Data management stack-based
3. VM-based
4. Hypervisor stack
5. SAN-based

1. Traditional

A traditional DR solution such as Microsoft Data Protection Manager, Veritas, or Commvault uses a client installed on each server to archive data in incremental backups to storage hosted either in the cloud or locally. In the event of a disaster, one of those backups can be used to rebuild new servers. These solutions are proven, do not require an environment to be virtualized, and typically do not require as much storage capacity as do other systems. They do, however, traditionally have high RTOs and RPOs, no application awareness, and little cross-platform support and are often unable to perform any source-side deduplication or continuous data protection (CDP).

2. Data management stack-based

Data management stack-based DR solutions such as Rubrik and Cohesity are relatively new to cloud disaster recovery. As the name implies, this type of solution replicates data at the data management stack level. As a result, this type of solution runs equally well with virtualized and non-virtualized machines. It can store data in an immutable format where massive changes can be verified, and malware caught before sweeping changes to DR data are made. This type of solution benefits from competitive RPO/RTO times and CDP, but it is slightly more cumbersome to host in the cloud because of the special software required to handle DR data at the data management stack.

3. Virtual machine-based

One of the more common new-generation DR solutions is VM-based. These solutions leverage available virtual machine infrastructure such as VMware, Nutanix, and Microsoft's Hyper-V by replicating and failing up/back the virtual machine using clients installed on the local site's virtual machine physical hosts. Because virtual machines are easier to failover and failback, they traditionally have competitive RTO and RPO times. Because virtual machines are easier to replicate, they also support continuous data

protection. The portability of VMs also makes setup of DR orchestration and automation extremely easy. A significant drawback of this solution, however, is that all critical systems must be virtualized before they can be protected.

4. Hypervisor stack-based

The hypervisor stack-based solution is not a complete solution on its own but usually relies on components of a VM-based solution to run. As a result, this solution is not hypervisor-agnostic and depends heavily on the local system's VM infrastructure solution. Currently, the only real hypervisor stack-based solution is VMware Cloud.

5. SAN-based

As the name implies, a SAN-based cloud DR solution requires the local site to have a storage area network. Because it leverages the SAN, this solution is the only one with hardware-based source-side deduplication, which can significantly speed up the process. This is also not a complete solution, because no SAN-based cloud DR solutions with built-in orchestration and automation components currently exist. SAN-based DR is best used in tandem with other solutions to set up a cloud DR solution for data that resides on an existing local SAN.

The tables below (figure 5) compare popular commercially available replication solutions, examining design, features, and pros and cons of each.

		VM Required	One Click Orchestration	Hypervisor	Cloud Providers	Monitoring and Analytics?	SAN Dependant?
VM	Zert0	Y	Y	Hyper-V, VMWare	Azure, AWS, IBM	Y	N
	Veeam	N	?	Hyper-V, VMWare, Nutanix	Azure, AWS, IBM	Y	N
	Actifio	Y	N	Hyper-V, VMWare	Azure, AWS, IBM, Google	Y	N
	Azure Site Recovery	N	N	Hyper-V, VMWare	Azure	N	N
Data management Stack	Cohesity	N	Y	Vmware, Hyper-V, Nutanix	Azure, AWS, Google	Y	N
	Rubrik	N	N	Hyper-V, VMWare, Nutanix	Azure, AWS, Google	Y	N
Traditional	Veritas		N		Azure, AWS, IBM, Google		N
	Comvault	N	Additional software req	Hyper-V, VMWare	Azure, AWS, Google	Additional software req	N
	Data Protection Manager	N	N	Any	Azure	N	N
SAN Based	Ontap Select	N	N	N/A	Azure, AWS, Google	N	Y
Hypervisor Stack	VMWare Cloud	Y	N	VMWare	AWS, IBM	Additional software req	N

	Traditional		Data management Stack		VM Based		Hypervisor Stack Based		SAN Based	
	Vendors: Veritas, Commvault, MS Data Protection Manager		Vendors: Rubrik, Cohesity		Vendors: Zert0, Veeam, Actifio		Vendors: VMWare Cloud		Vendors: OnTap Select	
Detached from infrastructure	No		Yes		Yes		Yes		Yes	
Source Side Dedupe	No		Software		VM Level		VM Level		Hardware	
RTO and RPO	High		Low		Very Low		Low		Low	
Orchestration Effort Required	High		Medium		Low		High		High	
Continuous Data Protection(CDP	No		Yes		Yes		Yes		Yes	
	Pros	Cons	Pros	Cons	Pros	Cons	Pros	Cons	Pros	Cons
	Established Base Line	No application awareness	Works on physical and virtual servers	More effort required to orchestrate VM backups	Hypervisor agnostic	All services must be virtualized	Detached from infrastructure	Not A complete solution -storage, orchestration, and analytics still needed		Not a complete solution- server, orchestration, and analytics still needed.
	Less storage capacity requrired compared to snapshot based systems	No cross platform support	Stored in immutable format					Not hypervisor agnostic		Limited to brand of SAN mfg

Figure 5: Replication Comparison Matrix

3.1.6. Templates and Automation

Templates allow cloud infrastructure conversion into infrastructure as code (IaC) and from IaC back to working cloud infrastructure. As a result, templates can be used to rapidly back up, restore, and replicate cloud infrastructure (see figure 6). Automation goes hand in hand with templates to use templates, virtual machines, and replication to orchestrate the disaster recovery tasks. These include initializing the cloud infrastructure, failover to that cloud infrastructure, configuring systems to run on that infrastructure instead of locally, and starting up those machines during a disaster scenario, with only a few clicks.

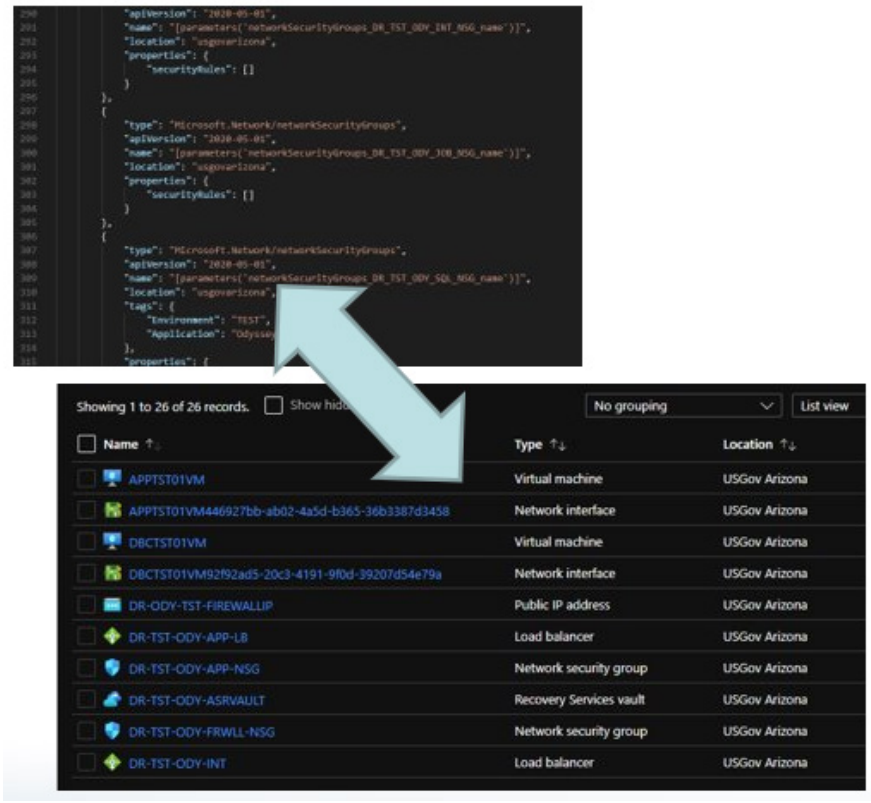

Figure 6: Sample Templates

Template and Automation Considerations

In evaluating automation, the number one priority is to verify that a given solution will work with a selected cloud provider. In this case, many cloud providers have built in tools that work well with their given solutions. In addition to the built-in automation tools in the cloud service provider's product offerings, a few open-source cloud infrastructure management solutions are available for automation. The most efficient way to manage cloud automation is with infrastructure as code. Some of these tools are orchestration tools; others are configuration management tools. Some tools are procedural; others are declarative. Above all, the selected solution must be compatible and optimized for managing cloud infrastructure for the selected cloud provider.

1. Orchestration or Configuration Management

Cloud orchestration tools look at the entire state of a cloud environment and attempt to ensure it is continuously in its desired state. If something is not functioning correctly, it automatically restores the system after reloading. Orchestration works well for environments that do not change often. Configuration management tools don't reset an entire system. Instead, they locally repair an individual issue. Some newer cloud management tools blur the line between these two types of tools.

2. Procedural and Declarative Tools

In choosing between procedural and declarative tools, just like with orchestration and configuration management, an overlap can sometimes exist within specific tools. A procedural tool is one that needs explicit direction and procedure defined in the code - and, it requires every small detail to be specified. A declarative tool 'declares' precisely what is needed and does not define the process of how to get that result. It takes a holistic view of the environment it is maintaining, and when the next check-up of that

environment runs, if there is any deviation to what was initially declared, it will attempt to restore the entire system to the originally declared state.

3. Cloud Provider Compatibility

Not all cloud infrastructure management solutions can provision all aspects of the specific cloud service provider's cloud infrastructure. It is important to verify which solutions work with the selected cloud provider before committing to one.

3.2. Cost Considerations

Cost is a significant consideration for cloud-based disaster recovery solutions. The following guidelines should be considered when factoring costs:

- Evaluate local investments in data centers and infrastructure against recurring cloud expenses. Cloud is expensive; however, as scale increases, so do the benefits of cloud DR versus on-prem.

- Optimize local resources. Right-size hardware and software licensing on the local systems before cost-estimating recovery solutions in the cloud. Many cloud DR solutions will replicate everything about a local server to the cloud. If a virtual machine is sized larger than needed locally for future growth, that machine will require more resources and, therefore, cost more when replicated to the cloud. Having a lean setup locally can potentially reduce costs.

- Use test cloud environments to prove concepts. Once a cloud infrastructure has been created, take advantage of this infrastructure to spin up a virtual environment for quick testing instead of using local resources. The effort to study cloud resource sizing and cost details would be well spent because considerable cost savings can be had if proper upfront research has been completed.

- Include licensing for the DR site. Virtual machines, network appliances, and applications running in the cloud all need licenses as well.

- Always-on cloud resources. Reserved Instances (Azure) for always-on cloud resources can reduce costs. Verifying which resources will work with a reserved instance is beneficial. However, it is important to minimize dependence on always-on resources whenever possible. Using IaC can assist in rapidly turning on and off cloud resources.

3.3. Current Feasibility of Cloud Disaster Recovery

Integration with any cloud environment increases complexity in many aspects of on-premise services. From configuring virtual hosts to work with cloud replication solutions, to setting up a VPN or high-bandwidth link to the chosen cloud provider, numerous additions to the on-premises site are required to implement a successful cloud DR solution. ***Therefore, at the time of this writing, cloud disaster recovery solutions are still more complex and costly than existing local site to site DR infrastructure.*** However, if applications are already hosted in the cloud, most cloud providers offer a straightforward way to setup disaster recovery replication from one geolocation to another, providing recoverability for cloud-based applications.

Setting up a new cloud presence requires staff or vendors/consultants who are up to date on their selected cloud provider offerings, DR solutions, virtualization, security, and cloud connection, and who possess expertise in how those technologies tie into the local site. It is important to realize the critical role that on-site technical expertise and resources play when working with a cloud presence.

For example, as cloud technology is continuously evolving and providers are continually optimizing their processes and procedures, resources must be committed to keep up with changes and training. What would be considered current at the time of setup, or go-live, may not be up to date in six months to a year.

4. Project Methodology

As previously mentioned, there are four recommended project phases that support a sound project methodology to define requirements, and select, design, and implement a modern disaster recovery solution. This section describes the detailed activities of each phase.

1. Phase 1: Discovery & assessment
2. Phase 2: Selection
3. Phase 3: Design
4. Phase 4: Implementation & pilot

4.1. Phase I: Discovery & Assessment

The Discovery & Assessment phase covers all tasks and activities related to the planning, hardware, policies, software, backup, and current disaster recovery for the local environment. The more work put into an assessment the better as this will greatly reduce the time required during design and implementation. If engaging a professional services vendor, being prepared for an assessment with documentation can streamline the experience and shorten the duration. The vendor will request a great deal of data to help with analysis along with performing their own environmental assessment to create a scope or plan.

4.1.1. Key Activities

The major activities in this phase are to:

1. Locate, create, and organize existing documentation
2. Review ITAC *Disaster Recovery* and *Next Generation Hosting Framework* documentation,
3. Define a scope and deliverables for the assessment
4. Identify tools to streamline assessment

This section provides details for each of these activities as well as additional considerations for the Discovery & Assessment phase.

1. Locate, create, and organize existing documentation. The data listed in this section and any other information collected on the environment will be useful during the assessment. All documentation should be revisited and updated with any changes throughout the subsequent phases of the DR project, including the following:

    - Data retention policies, current disaster recovery plans, backup and restore plan
    - Hardware, software, and licensing inventory, SaaS or cloud services
    - Network diagrams, IP address assignments, monitoring Infrastructure, security policies
    - Service catalog, service documentation

2. Review ITAC *Disaster Recovery* and *Next Generation Hosting Framework* documentation. The ITAC *Disaster Recovery Framework* is an in-depth document that provides guidance on DR maturity analysis, standard RTOs and RPOs, service prioritization, setting DR standards, providing reference and establishing methods of DR, and offers a build-your-own DR planning tool. The ITAC *Next-Generation Hosting Framework* provides tools and guidance on moving services to the cloud.

3. **Define a scope and deliverables for the assessment.** The scope of work for this phase of the project should focus on assessment and discovery by the court or a third-party vendor. A final assessment should create a report that documents the existing environment, readiness, and additional steps necessary to implement a cloud-based disaster recovery solution. The design and implementation assessment should be presented with solution design options available and define the scope of work for each option. Ensuring that the assessment phase is independent of any implementation processes is advantageous. The implementation should be clearly understood to be a subsequent phase meant to provide an opportunity to seek additional assessments, seek proposals from multiple vendors for the next engagement, or implement the solution with in-house resources. Discovery can be performed in-house by using tools and referencing whitepapers and knowledgebases. In this scenario it may negate the need for vendor assistance. The information gathered can be provided to a vendor for an assessment report, solution design or provide information to internal resources tasked with selecting available solutions.

4. **Identify tools to streamline assessment.** Many cloud service providers have developed calculators and tools that can be leveraged to streamline the Discovery & Assessment phase of a cloud disaster recovery project. These tools can be used by staff or vendors. Following are examples of available tools:

    - **AWS Cloud Adoption Readiness Tool(CART):** https://cart.splashthat.com/ is a survey designed to provide recommendations on migration to the AWS.
    - **Azure Migrate Service** https://aka.ms/azuremigrate/getstarted can assess Hyper-V and VMware server environments, databases, and web apps. Microsoft offers a virtual appliance that will collect metrics or complete a CSV template with server details. Azure Migrate Service also offers a dependency analysis at https://docs.microsoft.com/en-us/azure/migrate/concepts-dependency-visualization, which is an important data point when relocating services to a new location.
    - **Service Map**: https://docs.microsoft.com/en-us/azure/azure-monitor/insights/service-map offers detailed service mapping, which gathers information about all TCP connected processes joined to the service.

4.1.2. Additional Discovery & Assessment Considerations

1. **A disaster recovery plan and strategy are essential.** Having a disaster recovery plan and strategy in place with draft policies that align with the organization's Continuity of Operations Plan is advantageous. Once a solid disaster recovery plan is in place with a defined RTO and RPO, it will be much easier to select a modern solution, including cloud-based options that meet the needs identified.

    Application dependency also needs to be mapped as part of the court's technology inventory to determine which services are necessary to protect for failover in the event of a disaster.

2. **Backup and disaster recovery services are different**. Backup and disaster recovery services may overlap but are best treated as two separate services. Some disaster recovery solutions will also perform limited archival of data. However, a tried and true backup solution will better handle data archive and retrieval.

When evaluating options, existing disaster recovery and backup services should be considered to determine if they will complement the new service, as some DR and backup processes and applications can conflict with others.

3. **Thorough discovery upfront leads to accurate planning.** Having a current and comprehensive discovery/assessment of existing infrastructure will lead to efficient and more accurate planning for a disaster recovery solution. External expertise to augment your discovery phase provides for a second set of eyes from an outside source (e.g., vendor, consultant, or peer resources) and can help identify any overlooked challenges.

4. **Virtualization is key.** Virtualization is a major component and is foundational to most replication solutions. Nearly every major DR solution on the market works better with virtual servers. It is highly recommended to have an environment virtualized before attempting cloud replication.

4.2. Phase II: Selection

The Selection phase focuses on the tasks related to evaluating a solution that both fits the current disaster recovery needs and provides a path for the future. Many technologies are available for an organization to modernize their disaster recovery solution and/or leverage cloud disaster recovery solutions. Selecting and implementing the correct solutions that work together to create a functioning DR solution requires weighing requirements, geographic location, local environment, and the skills/abilities of the implementation team. If all these factors are considered prior to selection, a clear choice will become apparent.

4.2.1. Key Activities

The major activities in this phase are to:

1. Evaluate provider offerings
2. Choose a replication solution
3. Determine connectivity requirements
4. Evaluate automation tools

This section provides details for each of these activities, as well as additional considerations for the Selection phase.

1. **Evaluate provider offerings.** When considering a DR solution, it is essential to determine if the goal is to go straight to a cloud hosting solution or if the solution will be locally hosted for the time being. The selected cloud provider plays a key role in selecting many other components of a disaster recovery solution, as does the decision to go to the cloud or host locally. The replication solution, automation solution, and even decisions on connectivity depend on the specific cloud provider or the on-site hosting location.

2. **Choose a replication solution.** Once a cloud provider or local hosting has been selected, it is time to choose a replication solution. It will narrow down the number of replication solutions based on the compatibility with that specific cloud provider, or the ability to perform local site-to-site replication. From here, compatibility with the local environment should be evaluated. It doesn't matter what other benefits or features a replication solution has if it requires an organization to restructure their entire local environment to facilitate its use. The shortlist of replication solutions that meet the

local and cloud environment compatibility requirements makes comparison of each solution's RTO and RPOs against operational continuity and recovery goals easier. Additional factors such as costs, functionality, and configuration options should be considered in making the final selection for the replication tool.

3. **Determine connectivity requirements.** Several connectivity options are available to connect a local environment with a cloud provider, with varying degrees of cost. It is always ideal to buy the fastest, most robust connection available. However, cost is usually a consideration; therefore, selecting the proper connectivity option is an exercise in determining the minimum requirements for replication to function properly. This will involve calculating the total amount of data to be replicated, ascertaining the time window allowed to replicate this data, and using this information on a bandwidth calculator to determine the necessary bandwidth. It will be possible to compare that number to the cloud connection options available for the selected cloud provider, while factoring in cost.

4. **Evaluate automation tools.** With a cloud provider, replication solution, and connection tying everything together in place, the next component in the process includes evaluating automation tools. Working through the sequence with the filters mentioned above should distill a shortlist of automation tools for selection. These tools can then be evaluated against business requirements such as activation of failover and failback services with predefined triggers and the ability to manage local infrastructure. It is also important to consider the expertise of the technology team and whether they have any experience with available solutions.

4.2.2. Additional Selection Considerations

1. **Built-in orchestration tools streamline deployment**. When evaluating solutions, those with built-in orchestration tools can reduce time needed for configuration and deployment. This allows for one less application to integrate with the cloud DR solution and can free up time otherwise used to research and test potential applications to work on other aspects of the project.

2. **Existing skillsets can simplify adoption.** For example, if staff are knowledgeable of existing Microsoft technologies, it may be beneficial to select MS Azure as a cloud service provider. Many aspects of integrating Azure to other Microsoft services are a matter of point-and-click configuration. And using an end-to-end Microsoft solution may result in simplifying the disaster recovery configuration process.

4.3. Phase III: Design

Once a solution has been selected, this next phase focuses on the tasks related to solution design.

4.3.1. Key Activities

The major components in the Design phase are:

1. Cloud infrastructure
2. Cloud connection
3. Replication
4. Automation

This section provides details for each of these components as well as additional considerations for the Design phase.

1. Cloud infrastructure. The first step in solution design should be the cloud infrastructure on the selected cloud provider. This infrastructure provides the skeleton on which the rest of the solution will reside. A proper cloud infrastructure design will consider security best practices while at the same time facilitating the necessary connections between services and the local site. It will provide a home for Priority 1 (P1) services during a disaster and be configured to allow future growth of additional services or test environments. Engaging with outside expertise during the Design phase is beneficial to assist in maximizing cloud connection options to ensure the chosen solution is sized appropriately to the project.

2. Cloud connection. Once the cloud infrastructure is in place, designing the cloud connection between the local site and the cloud provider is possible. Replication, planned failovers, and failbacks will take place across this connection, so it is important to size it correctly. Devices from the local site must be able to communicate to DR storage in the cloud provider.

3. Replication. The replication solution is key in cloud-based disaster recovery. It controls what local services replicate to the cloud, how often, and major portions of how a disaster failover will occur. It is important to keep in mind which servers are Priority 1 (P1), including what the servers are hosted on, and how they will replicate to cloud storage. Many solutions will replicate across a public internet connection by default. However, it is also possible to set them up to use a dedicated private connection instead. There are often some smaller details of failover and failback that the solution is unable to manage.

4. Automation. The smaller yet oftentimes crucial failover/failback details not handled by the replication solution can be covered with a cloud automation solution. During the Design phase, it is important to identify what aspects of the failover will *not* be covered when clicking "failover" on the cloud replication solution. Oftentimes, automation needs to cover things like DNS, virtual machine startup order, and in some cases, run specific scripts to prepare a newly failed-over server for the cloud.

4.3.2. Key Activities

1. Predetermine naming conventions. Use naming conventions along with any grouping and tagging options available. Many cloud providers do not allow resource names to be changed or make it very difficult to change them once they have been created in the cloud. Combine this with the need to use an accurate naming scheme to quickly identify a specific resource out of a large list of every cloud resource in operation, and a solid naming scheme becomes extremely important.

2. Test new ideas using cloud infrastructure. Cloud offers the ability to take advantage of quick spin-up/spin-down of services to rapidly deploy test environments and potentially save on local infrastructure investments. Once the cloud infrastructure is ready, it becomes possible to test new ideas using cloud services quickly.

4.4. Phase IV: Implementation & Pilot

Once a solution design is complete, this next phase focuses on the tasks related to implementing and piloting a solution.

4.4.1. Key Activities

The step-by-step activities of the Implementation & Pilot phase are to:

1. Set up cloud hosting
2. Assign appropriate management roles to architects and administrators
3. Configure cloud IP range
4. Build necessary VNets or VPCs and add needed subnets
5. Set up and apply security groups to secure subnets
6. Configure peering of VNets or VPCs
7. Connect cloud environment to local site
8. Add necessary jump boxes, Active Directory, and any necessary DNS or DHCP resources
9. If required, set up a firewall and public IP addresses
10. Prepare the source environment
11. Deploy, configure, and verify the replication solution
12. Configure high-speed cloud connections to local site
13. Initiate pilot replication
14. Pilot failover exercise
15. Run pilot services in the cloud
16. Pilot failback exercise

This section provides details for each of these steps as well as additional considerations for the Implementation & Pilot phase.

1. Set up cloud hosting. Depending on the existing use of a public cloud, it may be necessary to enter into a cloud hosting agreement as a first step. Public cloud providers bill monthly based on consumption. The organization's finance team should be prepared to process invoices for subscription services. During the initial setup of a cloud provider, segregating the disaster recovery-related resources into a separate subscription or organizational unit is recommended, so that costs and security can be isolated. As the cloud solution is built out, further resource groups will be needed.

2. Assign appropriate management roles to architects and administrators. Prior to configuring anything in the cloud, it is important to lay out security and management roles early to keep rights and roles organized while allowing key team members access to assist with cloud infrastructure creation and future management.

3. Configure cloud IP range. Before creating any VNets, VPCs, or subnets, an IP range that will be used on the cloud provider must first be identified. Often services such as Amazon or Azure will provide a default range; however, if a hybrid cloud approach is desired, a range must be selected that will work with the local site.

4. Build necessary VNets or VPCs and add needed subnets. At this point, it is time to begin laying out the necessary infrastructure. If the necessary steps were taken during the Design phase, each resource should be created to that design's specifications. Take

extra care to use an easily decipherable and scalable naming convention on all new resources as many cloud providers don't handle name changes well.

5. Set up and apply security groups to secure subnets. Performing this step as soon as possible is key to quickly securing a presence in the cloud. It does not have to occur immediately after setting up VNets or VPCs and subnets; however, it should be performed before any external connection is made to the cloud site.

6. Configure peering of VNets or VPCs. This allows resources or virtual machines to communicate with other VNets or VPCs. It is recommended to set this up early to allow for testing of any jump boxes, firewalls, or virtual machines.

7. Connect cloud environment to local site. Be aware that many dedicated options like Azure ExpressRoute or AWS Direct Connect require additional lead time from the provider and cloud expertise to create those connections, and once they are in place, they can be costly. It is recommended to setup a bare-bones VPN connection during setup for a quick, cheap, alternative.

   Once this step is complete, the local site and cloud site should be able to communicate with each other.

8. Add necessary jump boxes, Active Directory, and any necessary DNS or DHCP resources. Once the cloud site can talk to the local site, it is then possible to set up and configure any necessary infrastructure services with cloud resources that will require an active connection to the local site.

9. If required, set up a firewall and public IP addresses. With cloud security, cloud networking, and infrastructure services up and running, it is now possible to set up and configure any public IP addresses and their associated firewall. It is advantageous to do this after configuring other virtual machine resources such as jump boxes or Active Directory servers, as they can assist in testing firewall connectivity while also providing authentication.

10. Prepare the source environment. Before a replication solution is deployed, the source environment should be prepped. At this point, virtual machines and virtual machine hosts should be organized. Many replication solutions benefit from organizing all production and all test VMs together on their own hosts, if possible. It is also advisable to make sure any on-site VMs are correctly sized, named, and configured before their initial replication to the cloud.

11. Deploy, configure, and verify the replication solution. Once the foundation has been created, deployment of the replication solution can begin. Depending on the solution, either the cloud side or the local site side needs to be set up first. Configuration of the solution should follow the solution design from the previous phase. It is recommended to initially test the solution on testing or staging servers. If possible, that testing or staging environment should be as close to production in size as possible to simulate similar RTO and RPO numbers during testing.

12. Configure high-speed cloud connections to local site. If setting up something more robust than the initial VPN connection, it is now possible to do so. Allow extra lead time if working with multiple providers to secure a dedicated line. Allow additional time for

a learning curve to set up and configure the connection on the cloud provider, or contract out to a knowledgeable vendor for assistance.

Once this step is complete, the environment is ready for replication.

13. Initiate pilot replication. Everything is now in place to begin replicating pilot data from the local site to cloud storage. Prior to performing a pilot failover, it is recommended to get a baseline RPO from several days of replication after the initial data sync.

14. Pilot failover exercise. After getting a baseline RPO and verifying RPO baselines, the next step is to perform a pilot failover of select services. It is advisable to time this step so a baseline RTO can be established. Be sure to involve subject matter experts to verify all failed-over services are working as expected.

15. Run pilot services in the cloud. This is optional but recommended. Once a service has failed over and been verified to work, continue operating it in the cloud for a short amount of time. During this time, monitor responsiveness of the service, and gather baseline information on its performance in the cloud. That data can be used in the future to verify cloud connection speeds, virtual machine sizing, and cloud infrastructure design.

16. Pilot failback exercise. Once all necessary cloud baseline data has been collected, the pilot service can now be failed back to the local site. During this failback, be sure to gather baseline data on how long this step takes and how much bandwidth is used. This information can be valuable when planning a real-world failback in the future. When all services have failed back, resume cloud replication to continue to protect those services.

    At the completion of this step, the disaster recovery exercise is complete.

4.4.2. Additional Implementation & Pilot Considerations

1. Backup cloud infrastructure as code. IaC should be backed up when possible to speed up deployments of new supporting infrastructure as recovery targets. Although virtual servers and network appliance configurations can be backed up using conventional methods, the time and effort put into designing and building cloud infrastructure cannot. This is where IaC plays a vital role, and it becomes possible to back up and replicate entire cloud environments with a few clicks.

2. There are benefits to a phased approach. A phased approach provides a learning curve for the implementation team, allows for validation of the solution, and builds confidence. It also allows a more managed one-at-a-time rollout of DR to the cloud on various services.

3. Vary test times to validate performance. RPOs can fluctuate with churn for some replication to technologies (ASR, Zerto, etc). With that in mind, it is advisable to test a new DR solution at various times and with different server loads.

4. Extending backups to the cloud. A backup strategy for cloud resources will need to be developed to protect items running on the cloud. Once an on-premises environment has been migrated up to the cloud, extending backups to those environments is essential.

5. Develop a cloud-based security strategy. A security strategy for cloud resources will need to be developed to protect items running on the cloud. Once an on-premises environment has been migrated up to the cloud, securing those environments in the same fashion as local systems is also important.

6. Combined or multifunction resources add complexity. Complexity will be added to cloud infrastructure, individual service failovers, cost breakdown/separation, etc. with combined or multifunction resources. In a perfect world, each service to be moved to the cloud would reside on its own virtual machine hosts, on its own virtual machines, have its own database, use its own storage and so on. In the real world, that is not always possible. Keep in mind that shared databases, storage, virtual machine hosts, etc. all add the complexity of determining what must fail up when attempting to configure and test one service at a time.

7. Failback considerations. The time it takes to failback (failback time objective or, FTO) is a major consideration as it can affect recovery exercises as well as actual incidents. It is one of the often-neglected details of cloud infrastructure. With all failbacks, there will be some downtime to copy down and sync cloud data back to local hosts. A planned failback can take quite a long time for large systems and this should be considered when attempting a failover/failback DR exercise.

   Additionally, a cloud disaster recovery solution should allow access to resources as the failback occurs. Some cloud DR solutions attempt to keep systems available for as long as possible during failback. Those are preferable to others which do not.

Several phases, tasks, and considerations are necessary for the recommended project management methodology to implement a cloud-based disaster recovery solution, and this section was intended to provide a comprehensive overview of all that is involved. The following section provides some real-life examples of pilots that were implemented to inform the recommendations contained within this *Disaster Recovery to Cloud Roadmap*.

5. Implementation Examples

5.1. Superior Court of California, County of Monterey (ASR / Hyper-V)

Monterey Court's disaster recovery to the cloud implementation replicates Hyper-V virtual machines to Azure over a 1 Gbps data center link using Azure Site Recovery.

5.1.1. Local Site

- **Virtualized?:** Yes
- **Virtualization Solution (if any):** Microsoft Hyper-V
- **Existing Backup Solution:** Microsoft Data Protection Manager to cloud cold storage
- **Preexisting DR Solution (if any):** Site to site Hyper-V cluster replication
- **Diagram:** Example of P1 service virtualization and local site to site replication

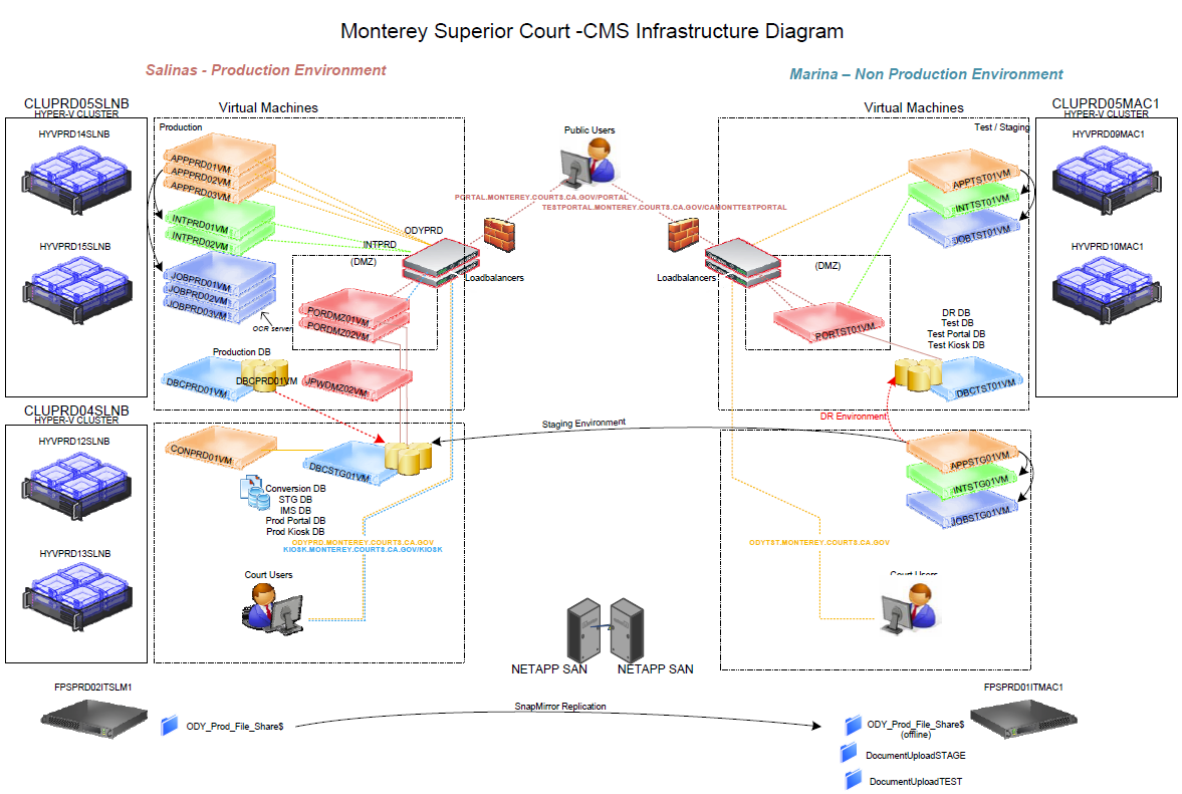

5.1.2. Connectivity

- **Existing Internet connection:** 1 Gbps AT&T
- **Dedicated cloud connection?**: Yes
- **Dedicated cloud connection provider:** 1 Gbps AT&T ASE to Equinix Datacenter. 1 Gbps cross connects to Azure inside datacenter
- **Diagram:**

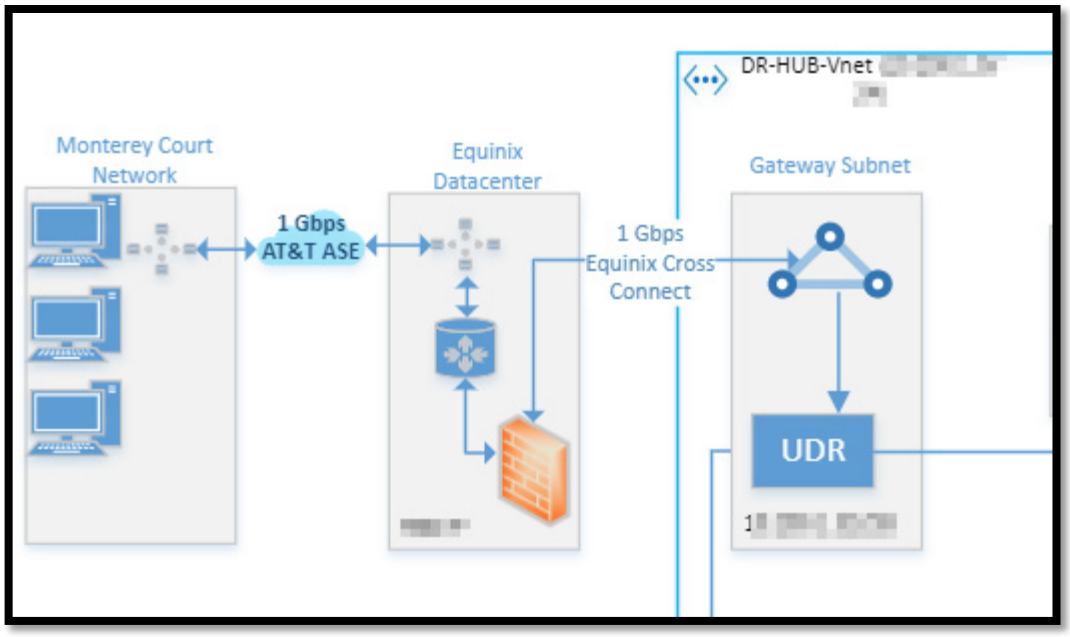

5.1.3. Connectivity Costs

- **Connectivity Monthly Cost:** $6,535.97/mo
- **Cost Breakdown:**

1000 Mbps Equinix Hosting VS 1000 Mbps AT&T Express Route

Item	Monthly
Equinix Datacenter Space / Power / Etc	$2,423.22
Azure Express Route Charges	$852.00
OPTIONAL: Internet Connection at Equinix	$2,000.00
AT&T ASE to Datacenter	$1,217.75
Emergency Redundant Network	$43.00
Total	**$6,535.97**

Item	Monthly
AT&T Redundant AVPN circuits	$12,073.30
AT&T NetBond, private peering	$3,969.00
Azure Express Route Charges	$852.00
Total	**$16,894.30**

5.1.4. Cloud Infrastructure

- **Infrastructure design:** Hub and spoke
- **Firewall:** Fortinet FortiGate
- **Public IP(s)?:** Yes, application specific
- **Additional security measures:** VNet Peering to segment applications, Network Security Groups for additional traffic control

- **Diagram:**

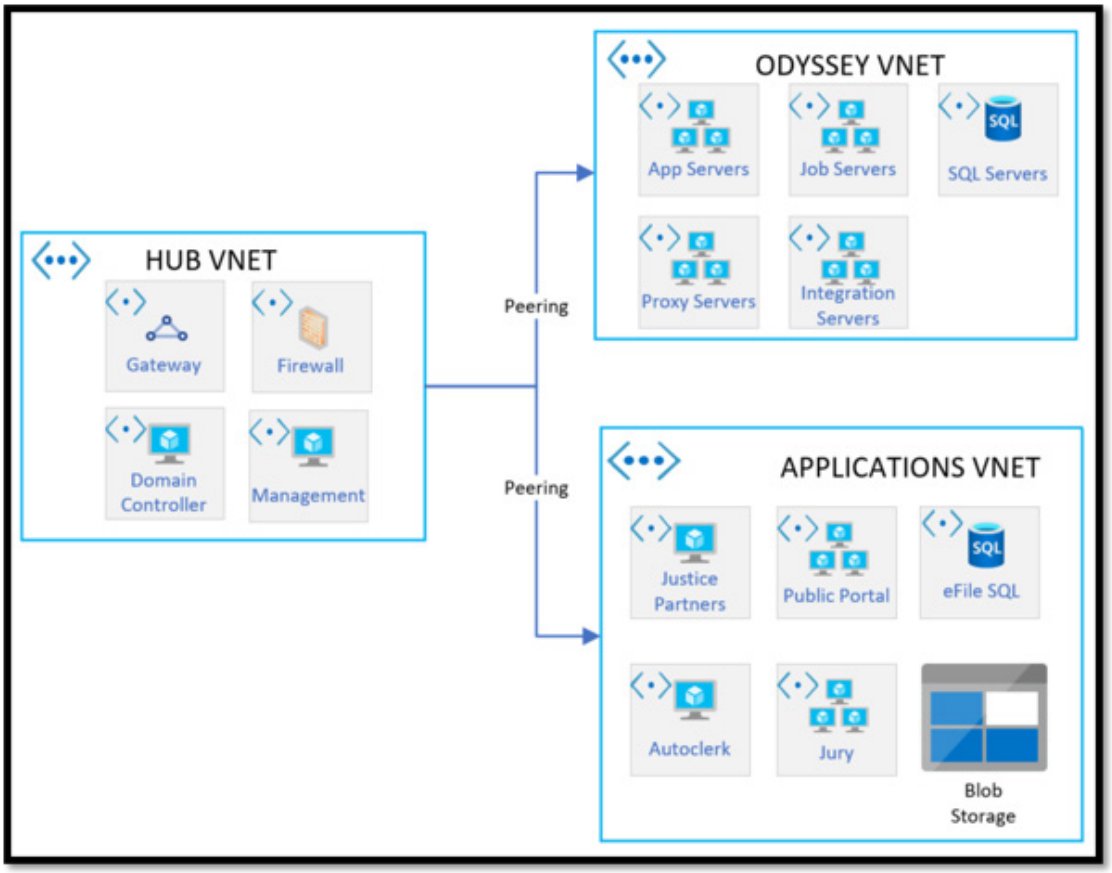

.

5.1.5. Storage

- **Storage Location(s):** Azure
- **Storage Type(s):** 1 Azure general purpose V1 Geo-redundant storage (GRS) per P1 application
- **Redundant?** Yes
- **If Redundant, what locations?:** Primary Azure USGov Arizona Region, Secondary USGov Texas Region
- **Storage size:** Currently 14TB for Odyssey Test site

- Diagram:

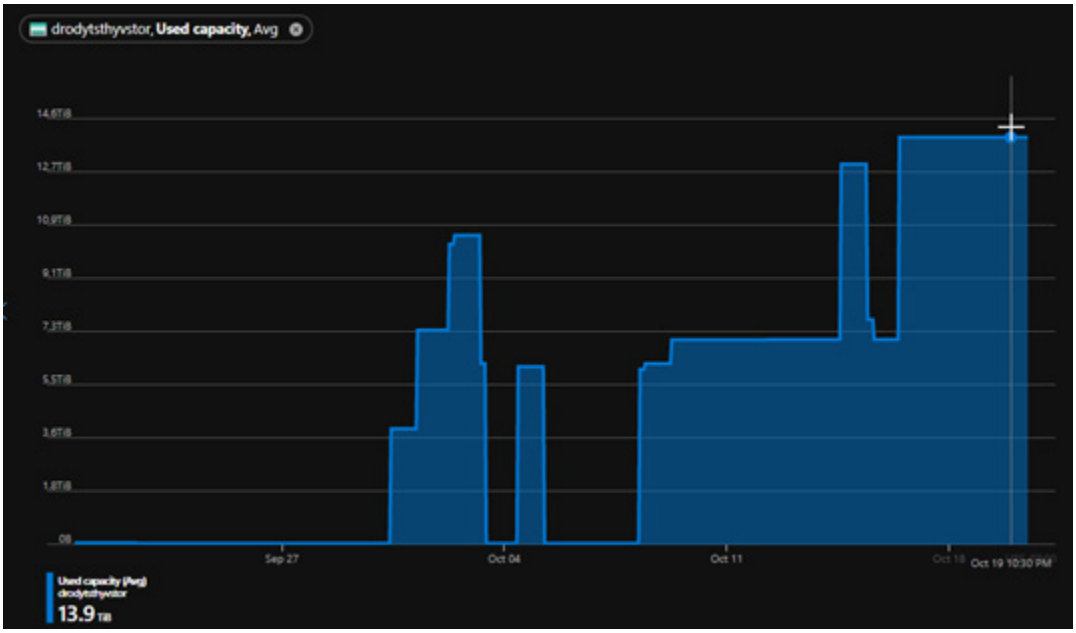

5.1.6. Cloud Storage and Hosting Costs

- **Monthly DR protection costs (please note if estimated):** $1,341.85 for Odyssey TEST environment (actual cost). Estimated approximately $100–$200 / month more for Odyssey Production.
- **Monthly live environment costs (please note if estimated):** $3,506.05 per month to host Odyssey TEST environment (estimated based on cost incurred from running TEST environment in Azure for 1 week).
- **Diagrams:**

Resource Name	Daily	Monthly
Odyssey Test Servers	$69.06	$1,036.50
Domain Controller	$3.58	$108.25
Blob Storage	$34.51	$1,035.30
ASR Replication Vault	$1.61	$48.30
Public IP address	$0.15	$4.50
Azure Firewall	$37.44	$1,123.20
ASR $25/mo fee per instance	-	$150.00
Running	**$146.35**	**$3,506.05**
Protected	**$39.70**	**$1,341.85**

5.1.7. Replication Solution

- **Cloud replication or site to site?:** Cloud
- **Replication solution:** Microsoft Azure Site Recovery
- **Estimated RPO:** 1-2 minutes
- **Estimated RTO:** <4 Hours
- **Diagram:**

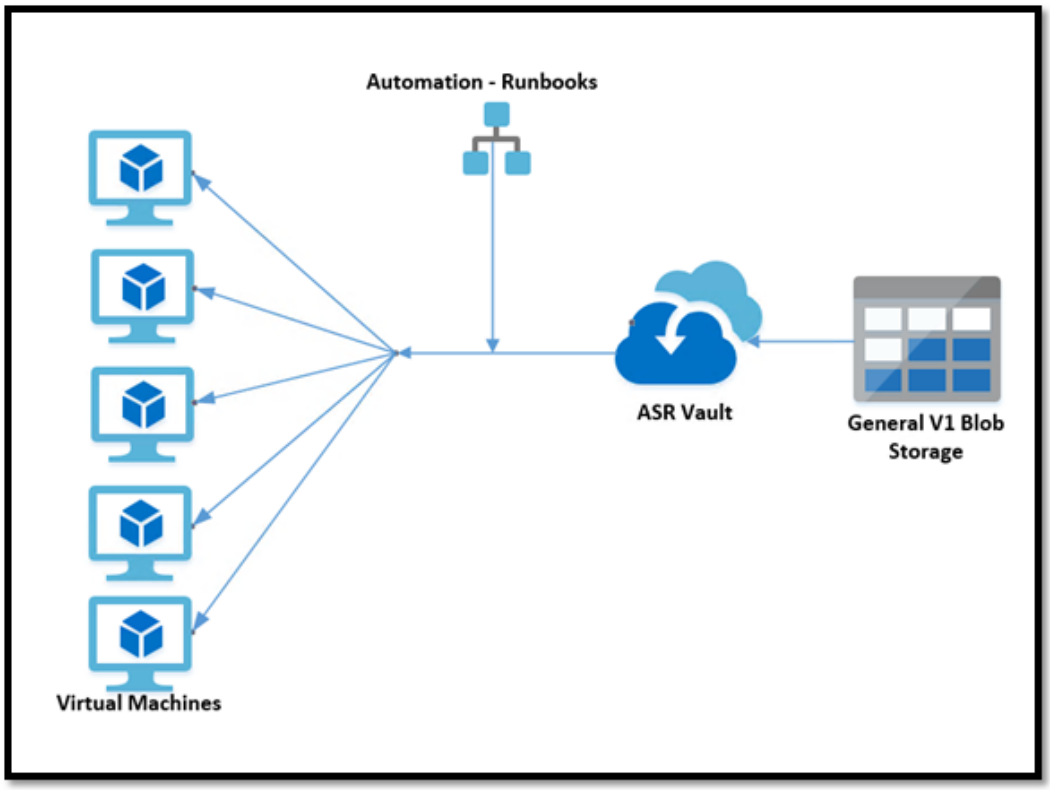

5.1.8. Automation

- **Failover Automation (if any):** ASR Recovery Plan utilizing Azure Runbooks
- **Infrastructure as Code (if any):** Microsoft Azure Resource Manager
- **Diagrams (if any):** See above
- **Infrastructure as Code Templates:** Can be obtained by emailing dr2crequest@monterey.courts.ca.gov ref: Brian Damschen

5.2. Superior Court of California, County of Santa Barbara (ASR / VMWare)

5.2.1. Local Site

- **Virtualized?:** Yes
- **Virtualization Solution(if any):** VMWare
- **Existing Backup Solution:** Veeam (local replicated to secondary datacenter
- **Preexisting DR Solution(if any):** Off-site backup to local secondary datacenter
- **Diagrams or graphics(if any):** None

5.2.2. Connectivity

- **Existing Internet connection:** 500Mbps L3,Wave
- **Dedicated cloud connection?**: No
- **Dedicated cloud connection provider:** NA
- **Diagrams (if any):**

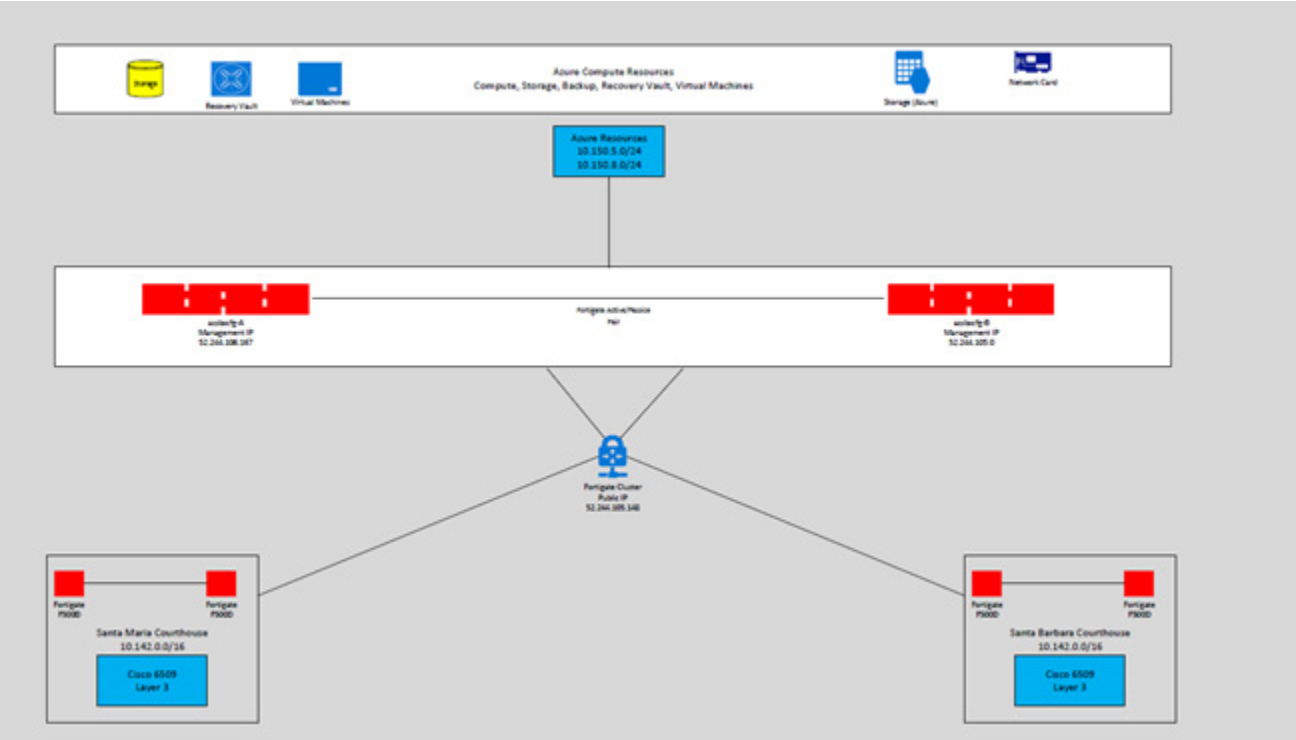

5.2.3. Connectivity Costs

- **Connectivity Monthly Cost:** $2000 per month, ISP costs
- **Please provide any cost data or diagrams you may have:**

5.2.4. Virtual Infrastructure

- **Infrastructure design:** Hub and Spoke
- **Firewall:** Fortigate
- **Public IP(s)?:** not currently

- **Additional security measures:**
- **Diagrams (if any):**

5.2.5. Storage

- **Storage Location(s):** Azure
- **Storage Type(s):** General purpose v1
- **Redundant?:** No
- **If Redundant, what locations?:**
- **Storage size:** 14TB
- **Diagrams (if any):**

5.2.6. Cloud Storage and Hosting Costs

- **Monthly DR protection costs (please note if estimated):** $2200
- **Diagrams (if any):**

			Values	
ServiceName	ServiceType	ResourceName	Sum of ConsumedQuantity	Sum of Price
Virtual Network	IP Addresses	Standard Static Public IP	4437	$23.52
	Peering	Egress	156.657103	$1.33
		Ingress	270.366278	$2.30
Virtual Network Total			4864.023381	$27.15
Azure Site Recovery	0	VM Replicated to Azure	14.08199048	$299.24
Azure Site Recovery Total			14.08199048	$299.24
Bandwidth	0	Data Transfer Out	42.470888	$3.94
Bandwidth Total			42.470888	$3.94
Azure Bastion	0	Basic	746	$150.92
Azure Bastion Total			746	$150.92
Grand Total			38954.20015	$2,239.20

5.2.7. Replication Solution

- **Cloud replication or site to site?**: Cloud
- **Replication solution:** Azure Site Recovery
- **Estimated RPO:** Every 60 minutes
- **Estimated RTO:** < 1 Hour
- **Diagrams (if any):**

5.2.8. Automation

- **Failover Automation (if any):** Not currently in place but looking into ASR Recovery plan using Azure runbooks.
- **Infrastructure as Code (if any):** None
- **Diagrams (if any):** None
- **Infrastructure as Code Templates:** None

5.3. Superior Court of California, County of Orange (Zerto / VMWare)

The Superior Court of Orange County's disaster recovery to the cloud implementation replicates VMware virtual machines to Azure over a 100 Mbps dedicated Microsoft Express route link using ZERTO.

5.3.1. Local Site (Irvine, CA)

- **Virtualized?:** Yes
- **Virtualization Solution(if any):** VMware
- **OS Stack (Windows):** Windows 2019, 2016, 2012 R2 and 2008 R2 (w/extended support)
- **OS Stack (Linux):** Oracle Linux 7.x, RedHat Linux 7.x
- **DB Stack:** MS SQL 2016/2012, Oracle 18c/12c
- **Existing Backup Solution:** Veeam local onsite backup and NetApp Snap Mirror to another Court location
- **Preexisting DR Solution(if any):** Limited using NetApp Snap Mirror array-based replication
- **Diagrams or graphics(if any):**

5.3.2. Azure Primary Site (US GOV Arizona)

- **Virtualized?** Yes
- **Virtualization Solution (if any):** Azure platform
- **Existing Backup Solution:** Azure backups
- **Preexisting DR Solution (if any):** Azure Site Recovery (ASR)

5.3.3. Connectivity

- **Existing Internet connection:** 250 Mbps AT&T
- **Dedicated cloud connection?** Yes
- **Dedicated cloud connection provider:** 100 Mbps redundant AT&T Netbond Express route connection to Azure Arizona.

5.3.4. New Connectivity (in progress)

- **Existing Internet connection:** 250 Mbps AT&T
- **Dedicated cloud connection?** Yes
- **Dedicated cloud connection provider:** 1000 Mbps redundant Megaport Express route connection dedicated for Azure Arizona (Primary) and Azure Texas (DR) and with both regions connected internally leveraging Azure global VNET peering.

5.3.5. Connectivity Costs

- **Connectivity Monthly Cost:** $5,078/month (will be greatly reduced as part of switching carrier from ATT to Megaport).
- **Cost Breakdown:**

Service	Qty	Monthly Recurring Charges
AVPN Circuit 100 Mbps Ethernet with Managed Router	1	$1,995.17
AVPN Circuit 100 Mbps Ethernet with Managed Router	1	$1,995.17
Netbond 100 Mbps Minimum Bandwidth Commitment - VNC (Private Peering, AZ)	1	$948.50
Microsoft Express Route, 100 Mbps metered	1	$140.00
Total Charges		**$5,078.84**

5.3.6. Azure Infrastructure (US GOV Arizona and US GOV Texas)

- **Infrastructure design:** Hub and Spoke
- **Firewall:** Palo Alto
- **Public IP (s)?** Yes
- **Additional security measures:** Network micro-segmentation, Azure user defined routes, (UDR's), Azure Network Security Groups (NSG's) for additional traffic control, Azure Network ACL's for storage accounts and Azure private link
- **Diagram:**

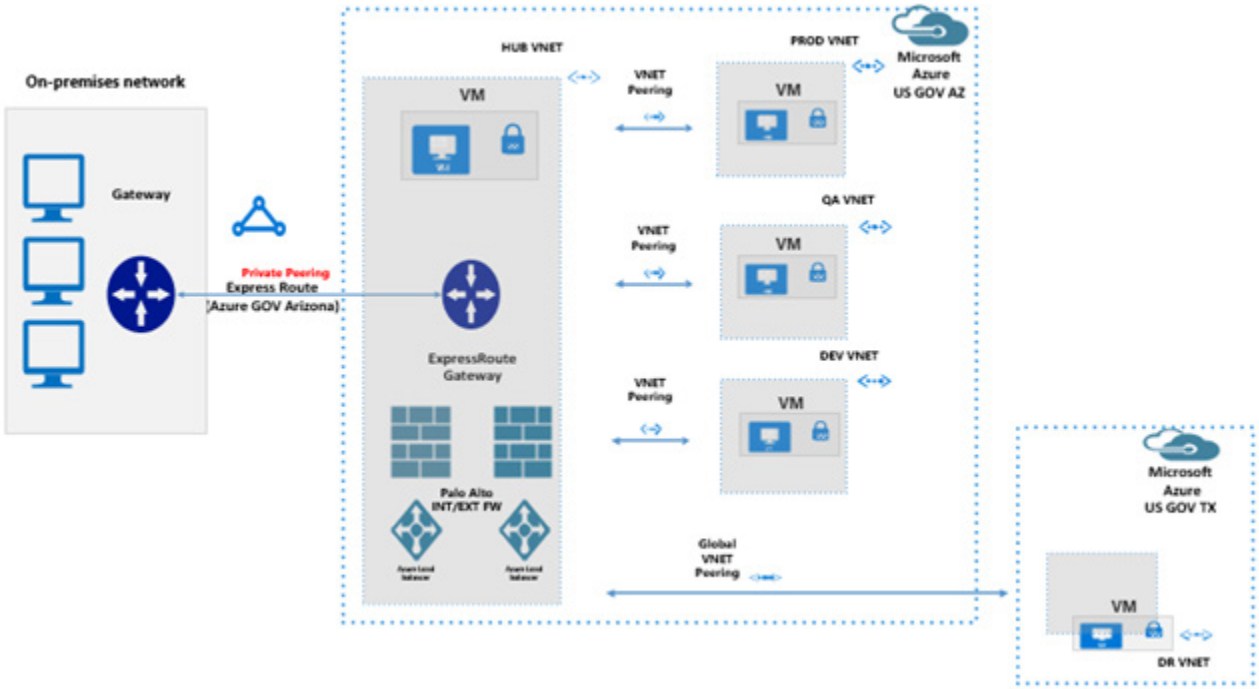

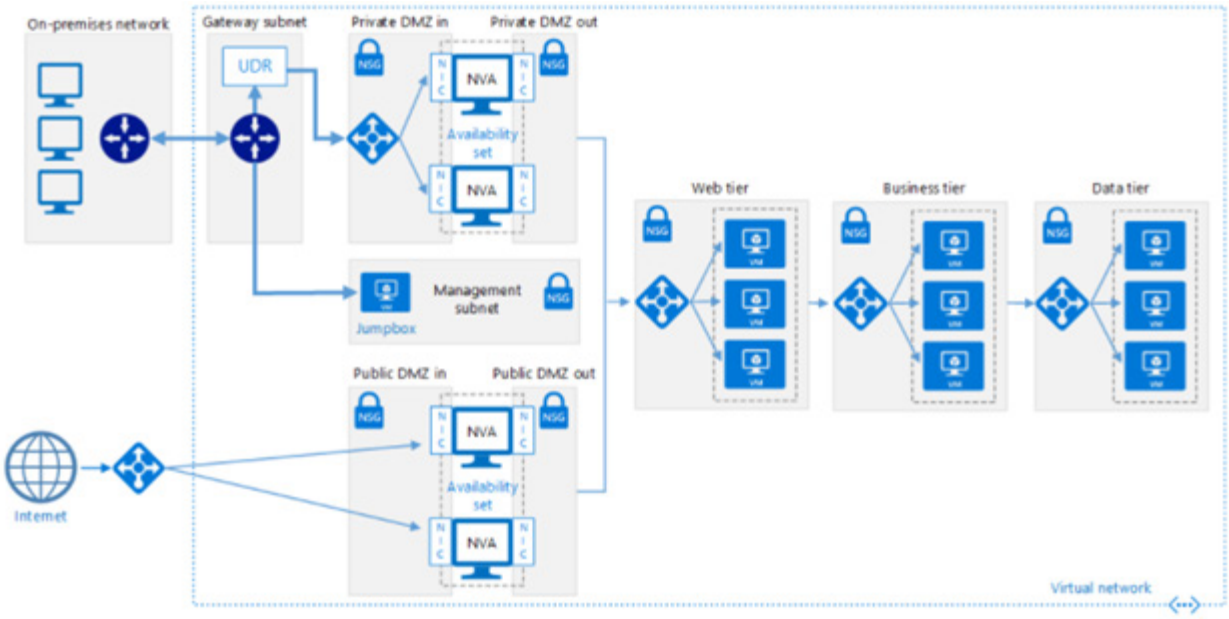

.

5.3.7. Storage

- **Storage Location(s):** Azure
- **Storage Type(s):** Azure general purpose v1 Local redundant storage (LRS)
- **Redundant?** No
- **If Redundant, what locations?** Primary US GOV Texas Region
- **Storage size:** Currently 60TB for 140 VM's (running on-premise and in Azure)
- **Diagram:**

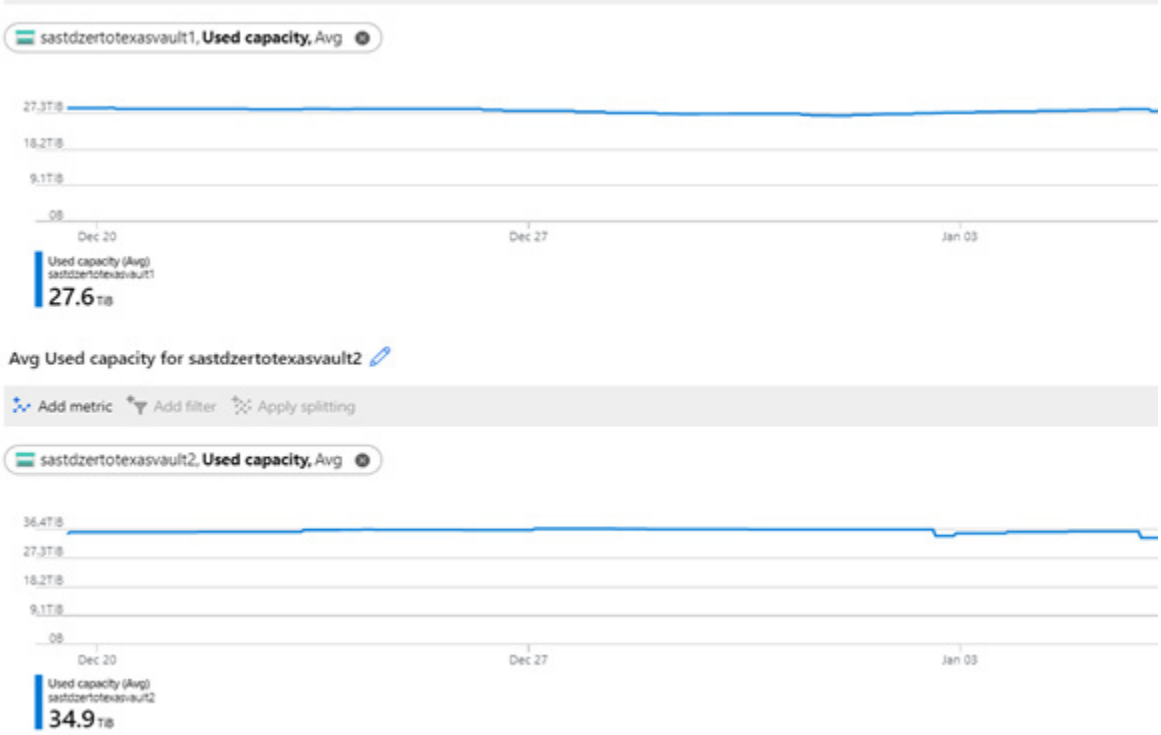

5.3.8. Cloud Storage and Hosting Costs

- **License Cost**: ZERTO Enterprise Cloud license and support
- **Monthly DR protection costs, for the protected workloads below (on-premise and in Azure)**

On-Premise (Irvine, CA)

- Tyler Odyssey PROD farm (Family Law and Juvenile CMS)
- Vision (Criminal CMS)
- Voyager PROD Farm (Civil, Probate and Mental health CMS)
- SharePoint PROD farm
- CRM PROD farm
- Public facing web apps
- IBM Data Cap PROD farm
- All Oracle and SQL PROD databases
- Internal web apps farm (IIS and WebLogic)

Azure (Arizona)

- Courts website
- Jury application
- Public facing web apps (IIS and WebLogic)

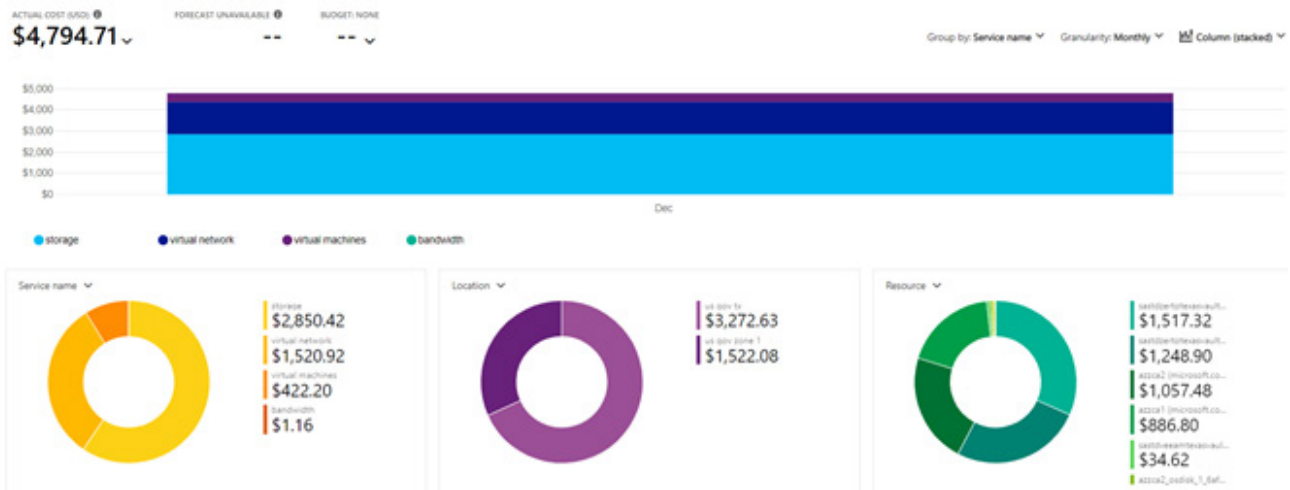

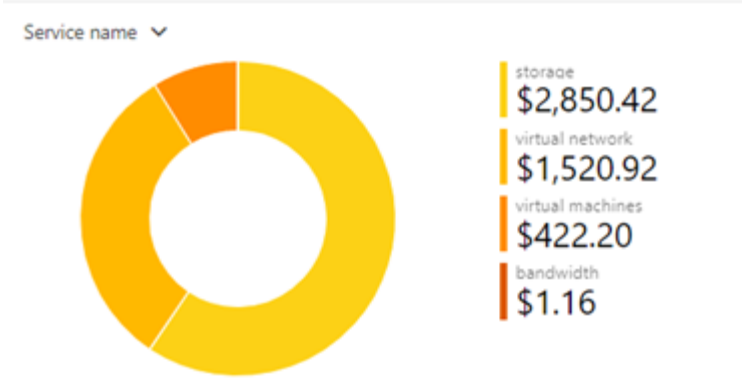

5.3.9. Replication Solution

- **Cloud replication or site to site?** Cloud
- **Replication solution:** ZERTO
- **Estimated RPO:** < 10 sec (VMware – Azure) and 1-2 minute (Azure-Azure)
- **Estimated RTO:** <4 Hours
- **Planning and Analytics:** ZERTO Planner
- **Diagram:**

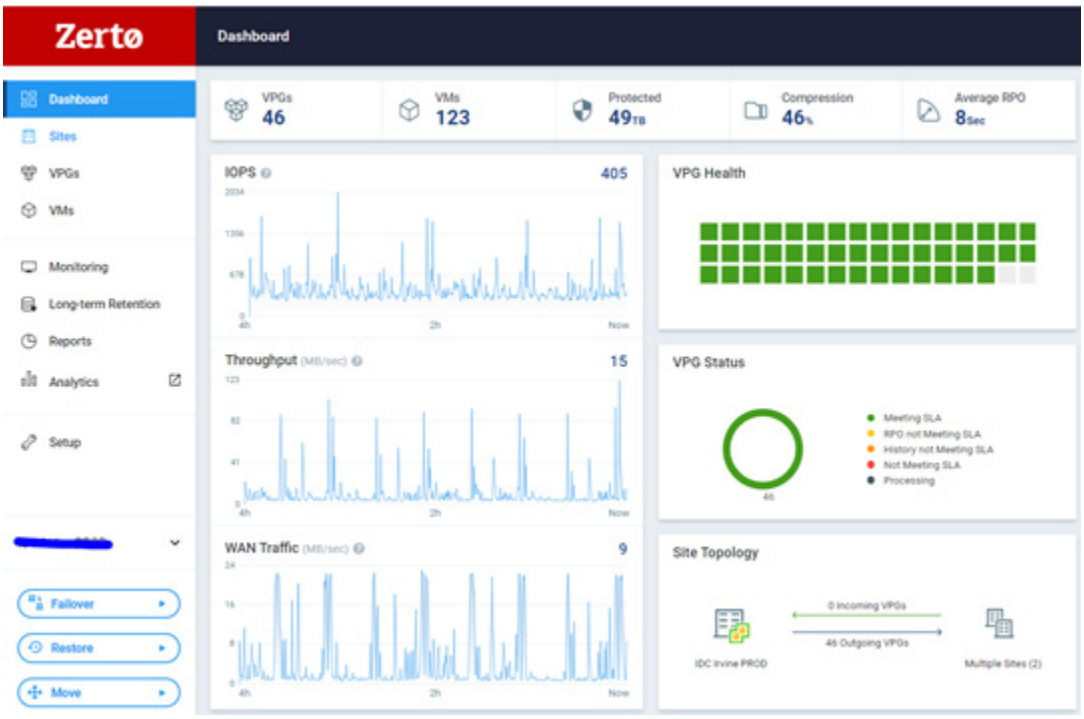

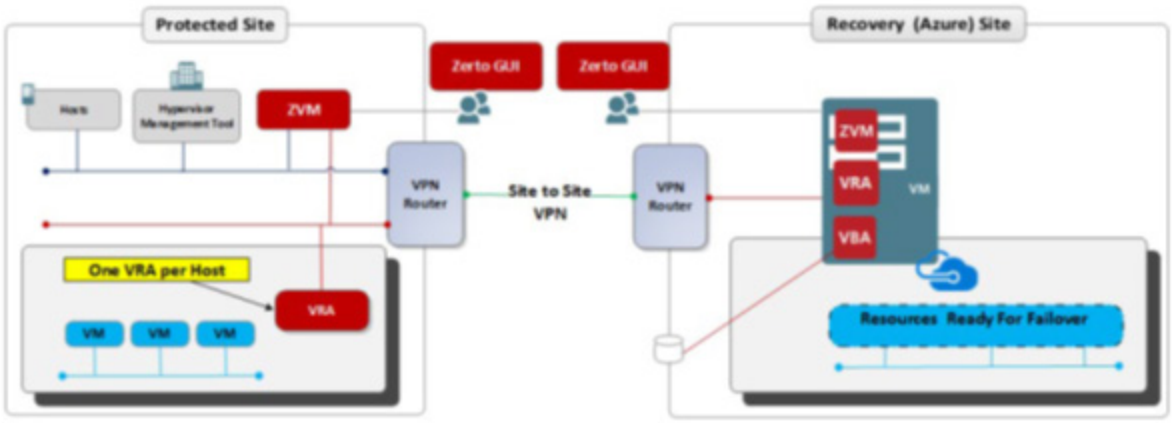

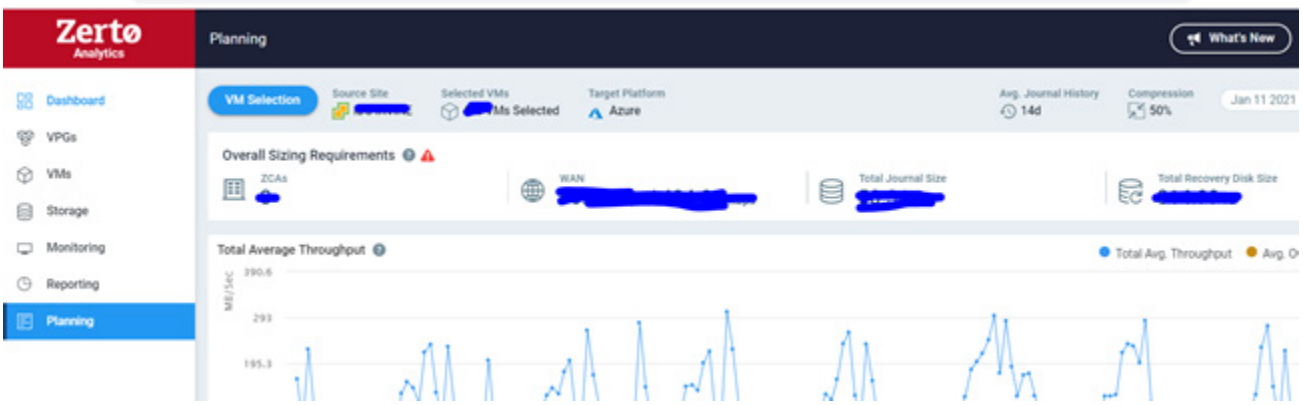

5.3.10. Automation

- **Failover Automation (if any):** ZERTO orchestration
- **Infrastructure as Code (if any):** Microsoft Azure Resource Manager
- **Diagrams (if any):** See above

5.4. Superior Court of California, County of Placer (Rubrik / VMWare)

The Superior Court of Placer County is utilizing a mirrored site to site 64TB Rubrik DR solution over a 1 Gbps connection.

5.4.1. Local Site

- **Virtualized?:** Yes
- **Virtualization Solution(if any):** VMWare
- **Existing Backup Solution:** Barracuda
- **Preexisting DR Solution(if any):** None
- **Diagrams or graphics(if any):**

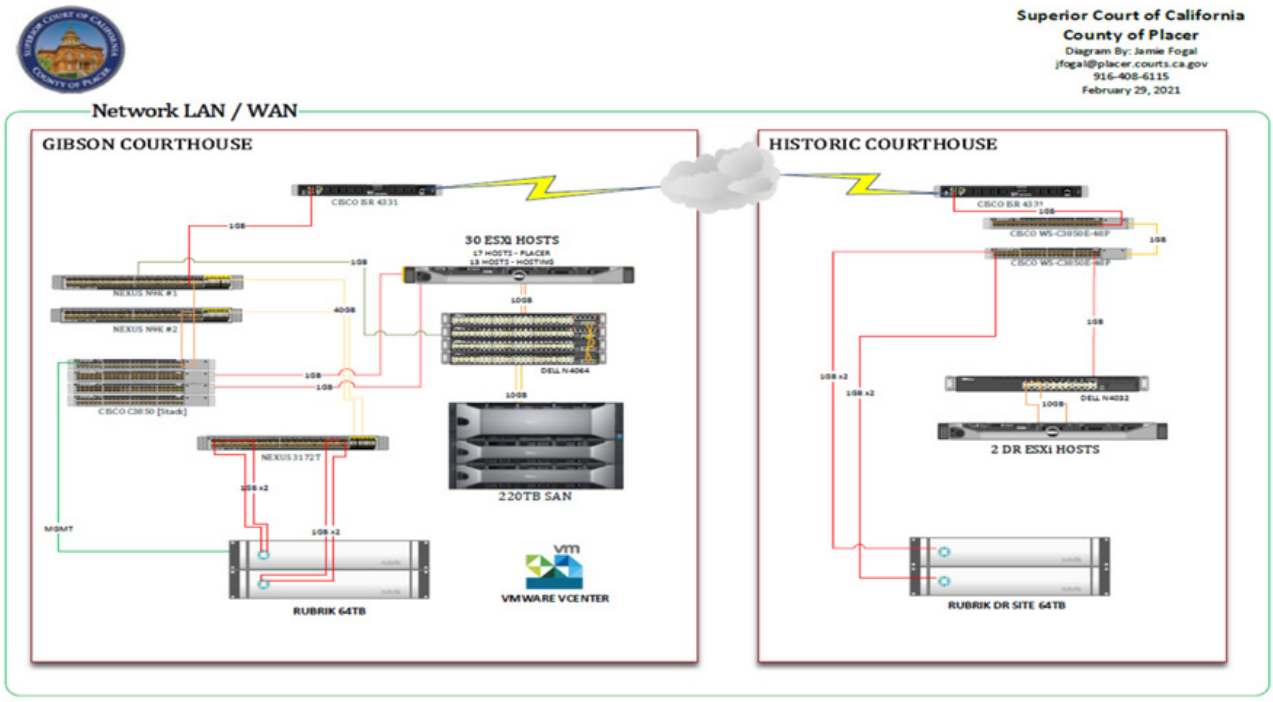

5.4.2. Connectivity

- **Existing Internet connection:** N/A
- **Dedicated cloud connection?:** No
- **Dedicated cloud connection provider:** No
- **Diagrams(if any):**

5.4.3. Connectivity Costs

- **Connectivity Monthly Cost:** $1000.00
- **Please provide any cost data or diagrams you may have:**

5.4.4. Storage

- **Storage Location(s):** Secondary Local Site
- **Storage Type(s):** Rubrik 64TB

- **Redundant?:** No
- **If Redundant, what locations?:**N/A
- **Storage size:** 64TB
- **Diagrams (if any):**

5.4.5. Replication Solution

- **Cloud replication or site to site?:** Site to site
- **Replication solution:** Rubrik
- **Estimated RPO:** 2-4 hours
- **Estimated RTO:** 4 hours
- **Diagrams (if any):**

5.4.6. Automation

- **Failover Automation (if any):** None
- **Infrastructure as Code (if any):**
- **Diagrams (if any):**
- **Infrastructure as Code Templates:**

Appendix A

Critical Court Services

Objective

This list of services identified as candidates for modern disaster recovery deployments to alternate datacenters, or to the public cloud, was derived from previous ITAC documents: *Disaster Recovery Framework*, *Next Generation Hosting Framework*, recommendations from the Disaster Recovery Phase I workgroup and included in multiple disaster recovery deployments in California courts of varying size. These are common services among all courts.

- Case Management System
- Electronic Filing
- Justice Partners Portal
- Case Data Portal
- Jury Management System
- Telephone System
- File Storage
- Network and Infrastructure to support local and hybrid connectivity