DEPARTMENT OF
VETERANS AFFAIRS

Office of Information and Technology
Service Delivery and Engineering (SDE)

VA Enterprise Disaster Recovery Service Tiers and Technology Solutions Standards

Version 1.0
4 September 2012

Revision History

Date	Version	Description	Author
9/4/2012	1.0	This merges the two documents that completed reviewed as SEDR12-0578 (DR Service Tiers) and SEDR12-0630 (DR Technology Solutions).	The MITRE Corporation Corporate Disaster Recovery Analysis Study Contract Number: VA798A-11-P-0015 Task Order Number: VA118A-11-0178 VA Project Manager - Zanna Child MITRE Project Leader - Mano Malayanur

Table of Contents

Table of Figures

Table of Tables

1 INTRODUCTION

1.1 Purpose

This standards document lists the acceptable and recommended specifications for Disaster Recovery (DR) *service tiers*[1] and *technology solutions*[2] for VA enterprise data centers.

Standard DR service tier definitions are required to consistently prioritize service restoration needs throughout the enterprise. Building DR solutions is best accomplished by first conducting a comprehensive business impact analysis (BIA) to determine the restoration priority and recovery time targets of information technology (IT) services that support business processes. IT services are grouped into DR service tiers according to their target recovery times. These DR service tier definitions include recovery time and loss of data objectives. These definitions, based on existing best practices and VA BIA data[3], are designed to fit within the current and anticipated VA architecture and satisfy business requirements in the event of a disaster.

This standard set of DR technology solutions provides a framework for selecting technology components when designing a DR technology solution. Standard DR technology solutions are presented for each application DR classification group defined in this document. DR technologies are presented using a four-layer solution architecture along with the reasoning behind the recommendations. Classification groups are defined in terms of a set of computing platforms and the standard DR service tiers; a set of computing platforms are also defined in this document. To determine the DR technology solution that meets the requirements of a VA application, an application will be mapped into an application DR classification group. The standard solutions suit the needs of the wide variety of applications hosted by VA data centers, and support VA goals of optimizing resources and increasing interoperability. The standard set of technology solutions, based on industry best practices, enables consistent deployment and maintenance of cost effective DR solutions throughout the enterprise.

---

[1] The recommended DR service tiers standard completed the System Engineering Design Review (SEDR) process of Enterprise Systems Engineering (ESE). The SEDR12-0578 review resulted in reference document *"VA Enterprise Disaster Recovery Service Tiers Standard, Version 1.0."*, which is an extract from the VA Corporate Disaster Recovery Analysis Study reference document *"5.1.5 B: Analysis Report Documenting Proposed Service Tiers."*

[2] The recommended DR technology solutions standard completed the SEDR process. The SEDR12-0630 review resulted in reference document *"VA Enterprise Disaster Recovery Technology Solutions Standard, Version 1.0,"* which is an extract from the VA Corporate Disaster Recovery Analysis Study reference document "*5.1.9: Recommended DR Technology Solutions, Platform Definitions, and Solution Classifications.*"

[3] BIA data was analyzed from VA's Security Management and Reporting Tool (SMART) database. This data set included the maximum tolerable downtime (MTD) cited by business units at field sites across all VA regions for 18,000 IT services.

1.2 Scope

This standard applies to:

- VA production applications deployed in VA data centers and the processes that support production applications.

1.1 Assumptions

This document makes recommendations based upon certain underlying assumptions. Should the validity of these assumptions change, recommendations contained in this report should be carefully re-evaluated for continued applicability. The following assumptions were made in preparing this report:

1. The DR service tiers recommendation assumes that the BIA information provided by the business owners is accurate.
2. The recommended DR technology solutions are meant to support a larger implementation effort which takes the application into account by utilizing information that would be provided by a BIA. It is assumed the implementation effort will include requirements analysis for solution selection and evaluation, as well design and implementation phases to fit within the existing application support lifecycle.
3. It is assumed that tertiary computing facilities providing additional levels of availability are not required for disaster recovery purposes.
4. Testing is vital to the success of a DR solution. While information on DR testing is included in this document, it is assumed that solution testing will be planned for as part of the application DR design process.
5. The IT strategic direction of VA is assumed from the FDCCI, 'Cloud First' policy, and OIT's Release Architecture v1.21 document.

1.2 Federal Policies and Standards

Several Federal Information Processing Standards (FIPS) and National Institute of Standards and Technology (NIST) policies are relevant to VA in the context of this project: FIPS 199, FIPS 200, NIST SP 800-34, and NIST SP 800-53. Additionally, VA Handbooks 6500, 6500.5, and 6500.8 address how the 800-53 controls are applied and implemented for systems and applications developed for, or used by VA. Table 1 below lists these documents.

Table 1. Standards and Policies

Standard	Title
FIPS 199	Standards for Security Categorization of Federal Information and Information Systems
FIPS 200	Minimum Security Requirements for Federal Information and Information Systems
NIST SP 800-34	Contingency Planning Guide for Federal Information Systems
NIST SP 800-53	Recommended Security Controls for Federal Information Systems and Organizations
VA Handbook 6500	Information Security Program Handbook

VA Handbook 6500.5	Incorporating Security and Privacy Into the System Development Lifecycle
VA Handbook 6500.8	Information System Contingency Planning

The controls listed in the above standards and policy documents are especially relevant during system design and implementation phases, and for this reason they are further described in Appendix A.

2 RECOMMENDED SERVICE TIERS

2.1.1 Description

Table 2 summarizes the recommended set of DR service tiers. Based on industry guidance about limiting the number of tiers to keep the solution simple and clear, it was decided that no more than 4 tiers should be created (see Rationale in Section 2.1.2). The first tier, Premium, combines those IT services having maximum tolerable downtimes[4] (MTDs) in the immediate, 4 hours, and 8 hours timeframes with both RTO and RPO of 15 minutes. This tier corresponds to the Premium tier in the notional model derived from best practices.

MTDs of 12 and 24 hours are combined into a High tier with RTO of 12 hours and RPO of 2 hours. The 48 hours, 72 hours, 7 days, and 21 days MTD categories are condensed into a single Medium tier with RTO of 48 hours and RPO of 24 hours. The final tier, Basic, is a combination of the MTD timeframes of 30 days and above, with RTO of 30 days and RPO of 7 days.

Notice that the highest tier in the notional model—Elite—is not recommended and at present is a placeholder in the VA tier set. The Elite tier is shown as part of the VA DR service model because it may become necessary in the future as business needs change and the cost of technology to support such a tier decreases.

[4] VA defines maximum tolerable downtime (MTD) as the earliest timeframe that an outage would have major or catastrophic impacts to the business's mission (Department of Veterans Affairs, 2011, "*2011 Data Center Consolidation Plan and Progress Report*").

Table 2: Recommended DR Service Tier Set

Mapping of BIA Data to Proposed Service Tier (% of IT Services)			Service Tier	RTO	RPO
Not a recommended tier at this time			Elite	Minutes	Virtually no data loss
Immediate	16.3%	29%	**Premium**	15 Minutes	15 Minutes
4 Hours	5.7%				
8 Hours	7.3%				
12 Hours	5.5%	13%	**High**	12 Hours	2 Hours
24 Hours	7.5%				
48 Hours	5.2%	24%	**Medium**	48 Hours	24 Hours
72 Hours	5.1%				
7 Days	5.3%				
14 Days	2.2%				
21 Days	5.7%				
30 Days	1.4%	34%	**Basic**	30 Days	7 Days
> 30 Days	32.6%				

2.1.2 Rationale

Number of Tiers. A four-tiered system was the most common DR approach we found in our research into industry and agency practices (see Figure 3-1 in "Summary Report of Service Tiers"). When crafting a recovery solution, it is desirable to reduce costs and increase the chances of operational success by limiting complexity and, therefore, the number of tiers. However, it is also important to ensure that sufficient variety of recovery targets and cost points are available to meet the needs of IT service owners. A four-tiered system strikes a balance between having too much variety (complexity) and needlessly increasing the sophistication and cost of a DR solution for applications that may not need it.

Elite Tier as a Placeholder. Achieving virtually no data loss and full recovery in just minutes requires a very sophisticated and expensive technological solution of fully automated failover systems and three data centers, two of which are located within a few miles. Few institutions require such stringent recovery times (financial institutions and intelligence agencies are among those that do), and such targets are difficult to achieve. Even Google, an IT industry leader in terms of capacity and innovation, which purports to have such an elite DR tier[5], has failed at times to meet its recovery objectives[6]. At this

[5] Google, 2010, *Disaster Recovery by Google*. http://googleenterprise.blogspot.com/2010/03/disaster-recovery-by-google.html.

[6] Needleman, Rafe, 2011, *CNET*. http://news.cnet.com/8301-17939_109-20102953-2/google-docs-suffers-30-minute-outage/, accessed April 1, 2011.

time no compelling business need was discovered in VA that indicated the need for an Elite tier of DR service.

Even VistA, one of the VA's most mission-critical IT systems, does not require virtually zero data loss because of the way in which the systems have been architected. The VistA system architecture consists of three copies of VistA: two running at two separate data centers with cache replication in between, and a read-only copy running in the hospitals. In theory, if a data center fails, the VistA instance is switched over to the standby copy in the second data center with only a few minutes' data loss. While the failover to the standby data center is taking place, the read-only copy of VistA at the hospitals provides reasonably adequate functionality for the hospitals to carry out their mission. This architecture makes hospitals less vulnerable to data center failures and makes RTOs of a small number of hours tolerable to the hospitals. In fact, the Veterans Health Administration (VHA) has proposed a two-hour recovery within its proposed service level agreement[7].

Until the required technology becomes more affordable or the VA's business need is great enough to warrant the high cost of an Elite tier, we recommend that this tier remain notional.

Premium Tier. Analysis of VA BIA data showed that 16 percent of IT services required immediate recovery, indicating a clear need for a DR tier that supports immediate recovery. The VA defines "immediate" as 15 minutes, and the RTO and RPO targets have been set at that point. Fifteen minute targets are achievable with database replication and automated failover between 2 data centers. The drawback of including services with 4- and 8-hour MTDs is the number of potential applications in the top tier (nearly 30 percent) and how that could drive up overall DR costs.

High Tier. The High tier has an RTO of 12 hours and an RPO of 2 hours. This tier serves the needs of those services with MTDs of 12 and 24 hours. The targets were selected to provide an option between the near immediate target of the premium tier above and the 2 days target of the medium tier below. The High tier corresponds to a DR tier currently employed by the VA, which should enable the VA to leverage appropriate knowledge and architectures already in place.

Medium Tier. The 48-hour recovery window reflects a preference by VA stakeholders with whom we spoke for meeting a need for a mid-level recovery solution constructed from standard technologies, such as SAN replication, at a reasonable price point. RPO has been set at 24 hours to align with the common operational practice of taking daily backups.

Basic Tier. The driving force behind this tier is the third of IT services having MTDs of 30 days or greater. The tier has an RPO of 7 days, however, if sites are already storing backups off-site more frequently than once a week, they should consider continuing that practice. A 7 day RPO does not preclude a team from restoring the data in this tier from a daily backup, if a daily backup is available. Again, this should enable the VA to appropriate knowledge and architectures already in place. Because many applications in this tier do not have spare equipment set aside for DR purposes, additional measures will be

[7] Department of Veterans Affairs, 2012, "*Draft ESLA 2.11 ISCP and DR 17 Feb 2012.*"

necessary to achieve a 30-day recovery time. An enterprise-wide strategy will be necessary to ensure spare hardware capacity is taken into account when planning for recovery of applications in this tier. For applications that do not have spare equipment readily available during a disaster the VA should maintain a predetermined list of server requirements and identify a dedicated procurement specialist to purchase the equipment with preauthorized purchase orders/credit cards (with estimated costs). This proactive procurement process would need to be incorporated into the annual DR testing scenarios.

2.1.3 Impact to VA

We anticipate that should the VA adopt these tiers as a standard across the VA, business units would have clear and consistent expectations of what DR tiers are supported across VA.

Naturally, instituting these tiers across the VA's regional data centers and Corporate Data Center Operations (CDCO) will impact VA Central Office, every VA region, IT customers (VBA, VHA, Cemetery), the BCM office, and OIT. *Fundamentally, this is not an IT change, but a change in the way the organization does business as it relates to DR.*

Impacts include the following actions that will need to be performed:

- Joint governance function between BCM and OIT established
- Tier determination driven by mission and business criticality, not just cost or funding availability. For instance, consider placing FISMA High applications automatically into the Premium or High tier unless specifically waived
- Technical architectures revised
- Communications for rolling out the tiers planned across the VA
- All DR plans and information system contingency plans (ISCPs) updated
- Service level agreements and operational level agreements with vendor contractors renegotiated
- DR funding scheme established and approved
- VA processes revised (for example, how to integrate the service tiers into the PMAS ProPath process so that business and application owners know when to decide on a tier, when to design the DR technologies, and when to test the DR solutions developed)
- In some cases, roles and responsibilities redefined
- Training and testing plans revised.

This is a large-scale change and should be managed as such—not handled individually within organizational silos.

A change of this magnitude on an enterprise level requires a large amount of coordination between multiple parties, therefore it is recommended that these service tier changes (as well as the additional changes recommended in other Study deliverables) be managed as a VA project, complete with a project management office and strategic guidance from organizational change management professionals, in order to smooth the adoption and ensure a successful deployment.

2.2 Critical Success Factors—Making the Tiers Work

Governance. Because of the complex interdependencies between VA IT services, the VA's BCM and OIT divisions will need to collaboratively develop policies, procedures, and criteria for assigning interdependent applications to DR service tiers. As covered in detail in "Summary Report of Service Tier Research," there are two ways in which to deal with interdependent services having differing recovery targets: 1) place all supporting services into the same tier as the dependent application that has more stringent recovery target or 2) allow for "degraded" service of the higher criticality service. A governing body of some kind will need to adjudicate conflicts arising from application owners who may not wish to pay for a higher level of DR service to ensure an application which depends upon their service can function within *its* restoration target. The governing group should also monitor the progress of the service tier implementation and adjust tiers if necessary to ensure they are in fact meeting the needs of the business units. Every few years after the tiers have been rolled out, the tier set should be re-evaluated for continued alignment to business requirements and available technologies.

Funding and Costs. Closely related to the need for governance is the issue of funding. Will the VA apply a chargeback model to individual service owners in order to pay for all or part of its DR costs, or will the DR service offerings be rolled under an overarching enterprise budget category for DR? In either case, the selection of service tiers should be first driven by the criticality of the IT service in support of VA's mission, followed by a rational discussion between the VA business and IT on funding DR to the appropriate level of service. For instance, consider placing FISMA High applications automatically into the Premium or High tier unless specifically waived. If current funding does not support the required service tier as determined by appropriate analysis such as the BIA, a governance mechanism should be in place to address the alternatives for funding or implications

There may be costs associated with transitioning from current service tiers to the recommended service tiers. The costs will depend on the amount of change needed to transition to the new service tiers. Additional costs will also be incurred if new services tiers are adopted in a data center, for example, if the premium service tier is adopted.

As stated in the methodology section, this study has taken the existing service tier investments into consideration, together with other important factors, such as industry frameworks and best practices, government and commercial practices, and business requirements. Since a key driver of this analysis is to align current VA practices and requirements with business needs and industry best practices, less weight has been given to the existing VA DR service tiers. However, it remains an important issue to understand and articulate the cost implications for the data centers to transition to the recommended service tier standard. Mission-critical owners may in some cases find the need to re-architecture their applications in order to switch to the new Premium tier. It is expected that the VA will conduct further investigation as to the ultimate cost of bridging the technological and process gaps to ensure that the data centers are able to transition successfully to the new DR tiers. For those data centers that do not have an established set of DR service tiers, a transition roadmap will need to be crafted.

2 PLATFORM BACKGROUND

A platform comprises hardware and software computing components, including generic types of physical server hardware, hypervisor (if used), operating system (OS), application development tools, database systems, and other common systems used to develop and run applications. Table 3 below shows examples of platform components. Vendor-specific products listed are meant only as examples and are not to be construed as endorsements.

Table 3. Examples of Platform Components

Category	Description
Server Hardware	Defined by CPU Instruction set architectures (e.g., x86, SPARC, ARM, IBM Power processors)
Operating Systems	Software that manages computer hardware and supporting services. (e.g., Windows, Solaris Unix, AIX, z/OS)
Hypervisor	Hardware virtualization software (e.g., Citrix XenServer, VMware ESX, Microsoft Hyper-V)
Database Systems	A data collection combined with a management system (e.g., Oracle, IBM DB2, MySQL)

Platforms are used to classify VA applications and match them to a DR solution. This limits the scope of available platforms and requires that they align with VA initiatives and industry trends. The following criteria are relevant when considering platform standards:

- Criticality of usage within VA
- Commonality of usage within VA
- IT strategic direction, such as the FDCCI and Cloud First Policy
- Maintainability and system development lifecycle as defined by enterprise architecture
- Additional characteristics that impact DR solutions

2.1 Industry Research

NIST, part of the U.S. Department of Commerce, recognizes three platform types within their publication "Contingency Planning Guide for Federal Information Systems" (Swanson, et al. 2010). These platforms are:

- Client/server systems
- Telecommunications systems (LANs, WANs)
- Mainframe systems

NIST recognizes that these three categories have distinct requirements and considerations with regard to contingency planning. The platforms' distinguishing features are the characteristics of the data and the types of external dependencies. For example, in a client/server platform, data can be distributed between client and server, as opposed to in a mainframe platform where large amounts of data reside primarily on the mainframe. An example of dependency differences can be seen in how a client/server is

dependent on the telecommunications systems, and the telecommunications systems in turn are dependent on the owners of the geographic regions they cross through.

Market analysis provides another source of information on platforms. In 2011, International Data Corporation performed server technology market analyses (Scaramella, et al. 2011). Their findings are for the worldwide market in 2010, of which the United States represented 38.2% of customer revenue:

- Server hardware popularity as per server shipments:
  - x86: 97.4%
  - RISC: 2.1%
  - EPIC: 0.5% (referring to Itanium architecture)
  - CISC: <0.01%
- Operating system popularity as per server revenue:
  - Windows: 46.0%
  - Unix: 24.7%
  - Linux: 16.9 %
  - Z/OS: 8.1%
  - i5/OS: 0.7%

While these percentages are not specific to the federal market, they clearly indicate the acceptance of x86-based hardware and operating systems such as Windows, Unix, and Linux in the server market.

The prevalence of these components is reflected in the basic service offerings of several federal data centers. The Defense Information Systems Agency offers support for the following types of equipment as part of their core services (Defense Information Systems Agency, Computing Services Directorate 2011):

- IBM/Unisys Mainframes
- IBM Mainframe (running z/Linux)
- Servers (Windows or Unix)

Corporate Data Center Operations, an organization which provides services to VA, advertises experience with the following platform components (Corporate Data Center Operations 2010):

> "Windows Server, SUSE Linux Enterprise, FreeBSD Enterprise, Oracle 10+ Application Server and Database, Z/Linux on IBM Virtualization, SUN Solaris Server, VMware ESX Virtualization, Red Hat Enterprise, SQL Server 2005 Cluster, IBM Mainframe, HP-UX Server, and SUN Solaris Virtualization."

While there are no clear standards for platforms, the use and support of certain platform components, such as x86-based servers, mainframes, Windows, Unix, and Linux, can be used to support logical groupings of platform components.

2.2 Existing VA Standards and Platforms

Two VA data centers were analyzed as part of the VA DR Analysis project. Inputs were taken from stakeholder interviews, documents, and discussions with government and contractor subject matter experts. These data sources form the basis for the information presented in this section.

2.2.1 VA Standards

OIT has published several enterprise standards that are relevant to platforms. Table 4 below shows these documents.

Table 4. OIT Standards Documents

Title	Date	Author
VA Enterprise IT Infrastructure Standard Server Platform - Production V1.0	November 18, 2009	Enterprise Infrastructure Engineering
VA OI&T Virtualization Platforms Procurement Guidelines Version: 1.2	March 22, 2010	Enterprise Infrastructure Engineering
OIT Release Architecture V1.21	November 30, 2011	Service Delivery and Engineering

The document that describes enterprise standards for platforms is scoped to address x86-based server hardware and operating systems. Within this document, four server classes are recognized, which are shown in Table 5 below.

Table 5. VA x86 Server Classes

Class	Description
A	VM hosts, heavy workload transactional
B	Heavy workload application servers, heavy workload content delivery, light and typical workload transactional
C	Light and typical workload application servers, light and typical workload content delivery, web servers
D	This class is based on cloud computing environments and has yet to be defined

The requirements for each server class are defined in some detail, and those that are relevant to platform components are listed below. Note that the standards should be consulted for full details.

- **Processor (classes A,B,C)**: Type − x86_64 (e.g., Intel Xeon, AMD Opteron)
- **Operating System (classes A,B,C)**:
  - Type − Linux preferred based on federal requirement to use open source operating systems. Windows acceptable. All must be VA-approved operating systems and compatible with hardware.
  - Version − Current target version of Linux or Windows per VA Technical Reference Model (TRM).

The procurement guidelines for virtualization platforms state that class A servers should be used. However, there is no guideline regarding the hypervisor.

OIT's Release Architecture document outlines the current operating environment for VA data centers as well as the specifications for new IT systems. It emphasizes a movement towards commodity hardware, virtualization, and the use of cloud offerings. Table 6 below summarizes the platform specifications from this document. Note that the future specifications may use anticipated names of products.

Table 6. Release Architecture Specifications

Component	Future Specification	Current Specification
Operating Systems	Windows Server 8, Red Hat Enterprise Linux v6.1	Windows Server 2008 R2, Red Hat Enterprise Linux 5.7
Virtualization	VMware TBD, Windows Server 8 Hyper-V	VMware vSphere 5.0 (or equivalent Type 1 hypervisor), Microsoft Server 2008 R2 Hyper-V (or equivalent Type 2 hypervisor)
Cloud	TBD based on NIST definition and use of FedRAMP	
Physical Servers	Follow established VA Enterprise IT Infrastructure Standard Server Platform document	
Database Products		
• **Caché**	Intersystems Caché v2011	Intersystems Caché v2008.2 ad hoc 9526
• **Oracle**	TBD (Oracle 12g)	Oracle 11gR2
• **Microsoft SQL**	TBD (MS SQL 2012)	MS SQL 2008
• **MySQL**	TBD	N/A
VistA		
• **Operating System**	In development	Linux front end/back end
• **Database**	Intersystems Caché v2011	Intersystems Caché v2008.2 ad hoc 9526
• **Server Platform**	Follow established VA Enterprise IT Infrastructure Standard Server Platform document	Follow established VA Enterprise IT Infrastructure Standard Server Platform document
• **Storage Platform**	Enterprise Class Tier 1 array (exact platform TBD)	HP EVA 8400

VA is also under a directive from the Office of Management and Budget to follow the "Cloud First" policy (U.S. Chief Information Officer, Vivek Kundra 2010). The strategy described in this policy asks that agencies use commercial cloud technologies, private clouds and, where feasible, regional clouds with state and local governments. One area the "Cloud First" policy does not address, however, is

infrastructure standards. In fact, there are currently no cloud infrastructure standards that apply to VA. In lieu of a standard, NIST has stated the following with regard to cloud infrastructure:

> "A cloud infrastructure is the collection of hardware and software that enables the five essential characteristics of cloud computing. The cloud infrastructure can be viewed as containing both a physical layer and an abstraction layer. The physical layer consists of the hardware resources that are necessary to support the cloud services being provided, and typically includes server, storage and network components. The abstraction layer consists of the software deployed across the physical layer, which manifests the essential cloud characteristics. Conceptually the abstraction layer sits above the physical layer (Mell and Grance 2011)."

Although VA defines specifications for x86-based servers and VistA platforms, there is no clear standard for cloud infrastructure, nor does it appear one will be defined given the inherent abstraction of hardware used to provide cloud services.

2.2.2 Existing VA Platforms

Part of the data gathering included documenting a sample of existing VA platforms. This information, which is available in reference document, *"5.1.3: Inventory and Discovery",* shows that VA uses many of the same platform components identified previously and is consistent with the current Release Architecture specifications.

3 RECOMMENDED PLATFORM DEFINITIONS

Based on VA's needs, from a disaster recovery perspective, there are four computing platforms: Cloud Eligible, Common Systems, Mainframe, and VistA. These platforms are one of the classification criteria for matching applications to a DR solution.

3.1 Cloud Eligible

The term "Cloud Eligible" is used to describe platforms with characteristics that are appropriate for being hosted on a public cloud service. The Cloud Eligible platform utilizes the software and hardware specifications of the Common Systems platform, defined in Section 2.3.2, as well as criteria which measure an application's ability to move to a cloud infrastructure. This platform is defined to support VA's ability to leverage public cloud services and is the only platform that characterizes the application in addition to the supporting software and hardware.

Cloud Eligible characteristics include:

- Each component is virtualized or capable of being virtualized. This implies the following:
  - x86-based architecture
  - Use of a standard hypervisor
  - OS: Windows or Linux
- The readiness of the application and data for the cloud:
  - Security may need to shift from being implemented within the application/network layer from the platform level.
  - The database and operating system must be supported by the cloud vendor.
  - The amount of data involved needs to be considered. Transferring large amounts of data to the cloud may not be practical.
  - Cloud Eligible applications should not be tightly coupled to non-cloud eligible applications.
  - Critical and core business services should not be placed in the cloud.

Criteria that should be considered when determining the cloud eligibility of an application can be divided into two categories: 1) Common criteria and 2) Application criteria. The Common criteria apply equally to all applications and provide a guide to selecting a cloud service vendor. The Application criteria detail the requirements an application must meet to be considered Cloud Eligible[8].

[8] Further details on these criteria can be found in reference "5.1.7 B: Public Cloud DR Viability"

3.2 Common Systems

Table 7 below shows the standard for the Common Systems platform. It consists of computing systems based on x86, SPARC, or standard hardware, running Windows, Linux or Unix operating systems, and a standard hypervisor, if virtualized. This platform is identified due to the widespread use of Windows and Linux systems, and their increasing virtualization. The Common Systems platform provides the basis for the Cloud Eligible platform; however, additional criteria must be met in order for a system to be considered Cloud Eligible.

Table 7. Platform for Common Systems

Component Category	Recommendation
Hardware	X86-based
Hypervisor	VMware, HyperV, Linux-based (e.g., Xen)
OS	Windows, Linux
Database	Oracle, SQL Server, MySQL
Application Development Platform	TRM approved

3.3 Mainframe

The term “mainframe” is meant to indicate hardware used to support IBM mainframe operating systems, such as z/OS or z/Linux running on a z/VM hypervisor. The definition includes z/Linux virtual farms. This platform is identified in support of mainframe usage within VA and the functions mainframes support. Table 8 below shows the standard for Mainframe platforms.

Table 8. Platform for Mainframe

Component Category	Recommendation
Hardware	Mainframe
Hypervisor	z/VM
OS	z/OS, z/Linux
Database	CICS, IDMS
Application Development Platform	TRM approved

3.4 VistA

VistA is recognized as a platform type given its importance and use within VA. There are two main components to VistA; the front-end (VistA FE) which provides the web interface and application layer, and the back-end (VistA BE) which hosts the database. From a DR solution perspective, the key feature of VistA is the use of the Caché database. The presence of a Caché database is significant since Caché data replication provides unique challenges. For this reason, the VistA back-end is the focus of any VistA DR solution discussion in this document. VistA uses the platform architecture shown below in Table 9 and includes VistA systems deployed on Linux.

Table 9. VistA Platform Architecture

Component Category	Recommendation
Hardware	DEC Alpha, x86
OS	Open VMS, Linux
Database	Caché, VA Fileman
Application Development Platform	CPRS, Windows, RPC Broker, Kernel, VA Fileman

3.5 Alternate Platforms

Certain applications will use a platform that does not fall into one of the four standard platforms. One possible solution for these applications is to migrate them, if feasible, to one of the platforms described above. However, with the advent of newer technologies, standards will need to be modified to remain current with the needs of VA. The resolution of alternate platforms with regard to existing standards will ultimately require further exploration and is outside the immediate scope of this document.

4 DR SOLUTION BACKGROUND

Disaster recovery is the process, policies, and procedures that pertain to the recovery of a business's or organization's IT infrastructure after a disaster. A disaster can be a natural or man-made event that results in a severe impact to the working environment, people, or facilities such that service cannot be provided. The DR solutions described in this paper focus on technology solutions for the recovery of applications within VA data centers at alternate processing facilities and have been designed to match the classification groups defined in *"5.1.6 A&B: Classification Groups Research, Proposal, and Summary"*.

Regular testing of DR plans is critical to train staff and to assess and validate DR components. Without regular, rigorous, and realistic testing, DR plans and solutions are almost certain to fail in the event of a disaster. Appendix B contains information on best practices and federal guidelines in this area.

4.1 Key Concepts

A key concept is the difference between high availability and disaster recovery. From a practical point of view, the intent of high availability is to address routine component, sub-system or system failures, while the intent of disaster recovery is to address catastrophic failures or events that could take down an entire site. HA system design and implementation are intended to maximize the time that a user is able to utilize a system. This is often accomplished through the use of redundant components and software. HA solutions can be and are implemented as part of DR, but not all HA solutions are pertinent since DR scenarios have different and often unusual requirements, such as the need to failover computing infrastructure across a significant geographic distance. Therefore, this document does not address HA solutions that focus on system redundancy at a local level. On the other hand, HA solutions with built in geo-redundancy capabilities can be good candidates for DR solutions.

DR technical solutions can provide many features, but the key attributes that determine the characteristics of a DR solution are the Recovery Point Objective and Recovery Time Objective that can be supported. The RPO refers to the amount of data loss that can be tolerated. For example, an RPO of eight hours requires that data can be recovered from a point within eight hours before disaster occurred. The RTO refers to the amount of time it will take for service to be made available from the point that disaster was declared. For example, an RTO of three days requires that it will take no more than three days from the time the disaster occurred for the application or service to be available. Table 10 below shows the characteristics of DR solutions with regard to the RTO and RPO they are able to achieve.

Table 10. DR Solution Features

	Near Zero RTO/RPO	RTO/RPO in Hours	RTO/RPO in Days	Unknown time to recover
Cost	Most expensive	>>	>	Least Expensive
Failover	Fully automated	Semi-automated	Semi-automated	Manual
DR Site	Has environmental, network, and computing infrastructure and is processing workloads.	Has environmental, network, and computing infrastructure but needs manual intervention to be made active.	Environmental, network, and computing infrastructure requires configuration before being able to process workloads.	DR facility may not exist. Infrastructure requires procurement activity.

4.2 DR Solution Considerations

4.2.1 Cost

Cost is an important consideration for any design. For this reason, an overview of DR solution cost considerations is provided here.

The cost of a DR solution can be high due to purchase of additional facilities, hardware, software, and testing. A DR solution can become more expensive as the RTO and RPO times become more demanding. The decision of which DR solution to use should be driven by data gathered from the BIA and weighed against the cost of service unavailability. Solution costs can come from a number of areas:

- **Infrastructure** – Consists of hardware, software, facility/space, power/cooling, network, and bandwidth. Solutions with demanding RTO/RPO will require more upfront investment to ensure that DR resources are available when they are needed.
- **Licensing** – Licensing for products can be an additional cost in a DR infrastructure, depending on the vendor's licensing model.
- **Support** – Vendor/contractor services might need to provide more assistance during a DR scenario. For example, if the solution needs to meet a twenty-four hour RTO, then twenty-four hour support needs to be available.
- **Design and Installation** – Based on the availability level the application needs to meet, additional design requirements may be imposed that can impact all phases of the design lifecycle (see section 4.2.2 below for more details).
- **Maintenance** – Existing DR solution infrastructure must be maintained. Servers and network components that must be able to process workloads need to be kept up to date with software updates and available configurations in order to meet RTO/RPOs.
- **Training and Testing** – Staff need to be trained in the use and maintenance of the DR solution. Regular testing is also required to validate and verify the solution's functionality. In addition to testing facility costs, personnel hours can be a significant cost factor depending on the number

of people and time involved. Appendix B contains a more comprehensive discussion of DR testing.

4.2.2 Application Design

Application design becomes increasingly important when the RTO and RPO requirements are more demanding. For an application to support low RTO/RPO, it must be able to leverage the redundant resources that support this goal. For example, a geo-cluster or stretch cluster is a group of servers operating as if they were a single machine, even though they are distributed across an extended distance. An application will only benefit from running in this environment if it is cluster-aware, which would enable it to trigger failover of its processes to redundant servers in the event of a disaster. An application's requirement to work with a DR solution must be kept in mind throughout the phases of an application's lifecycle:

- Analysis – System analysis and requirements definition
- Design – Description of the desired features, processes, and documentation
- Implementation – Acceptance, installation, and deployment
- Testing – Verification and validation of the product
- Maintenance – Changes and upgrades

Each of these phases will be impacted by the application's need to meet the availability requirement, and each has the potential to increase the cost of the solution.

4.3 DR Solution Methodology

A companion deliverable to this document, *"5.1.6 A&B: Classification Groups Research, Proposal, and Summary"*, outlines a methodology for matching applications to DR technology solutions. Each application is classified into a group based on its RTO, RPO, and platform; a solution is then defined for each classification group. To avoid prescribing a specific solution that might not take into account all of an application's DR requirements, the following solution approach was used.

Each DR solution is described by a list of features and a set of technology components that can meet the RTO, RPO, and platform requirements of the specific classification group. Design considerations and factors that can impact the use of particular technologies have also been included. The result is that each DR solution provides a technology selection roadmap.

These solutions are organized using a four-layered model based on the three-tier software architecture, chosen for its applicability when organizing application-focused DR technologies. The underlying fourth tier, "Common Infrastructure," is needed to represent infrastructure such as facilities and network components. These are not DR-specific items, but are necessary to support a solution.

4.4 DR Technology Solution Detailed Information

Each DR solution provides a technology selection roadmap. Each solution is supported by detailed research into the DR features of the suggested technology components, which can be found in *"5.1.10 B: Detailed DR Solutions"*. Related document, *"5.1.10 A: Detailed Solution Documentation Formats"*, provides the introduction, methodology, and formatting for 5.1.10 B. Figure 1 below shows a logical

representation of the information relationship between this document, *"5.1.9: Recommended Platforms and DR Solutions"*, and 5.1.10B.

Figure 1. Relationship between 5.1.9 and 5.1.10

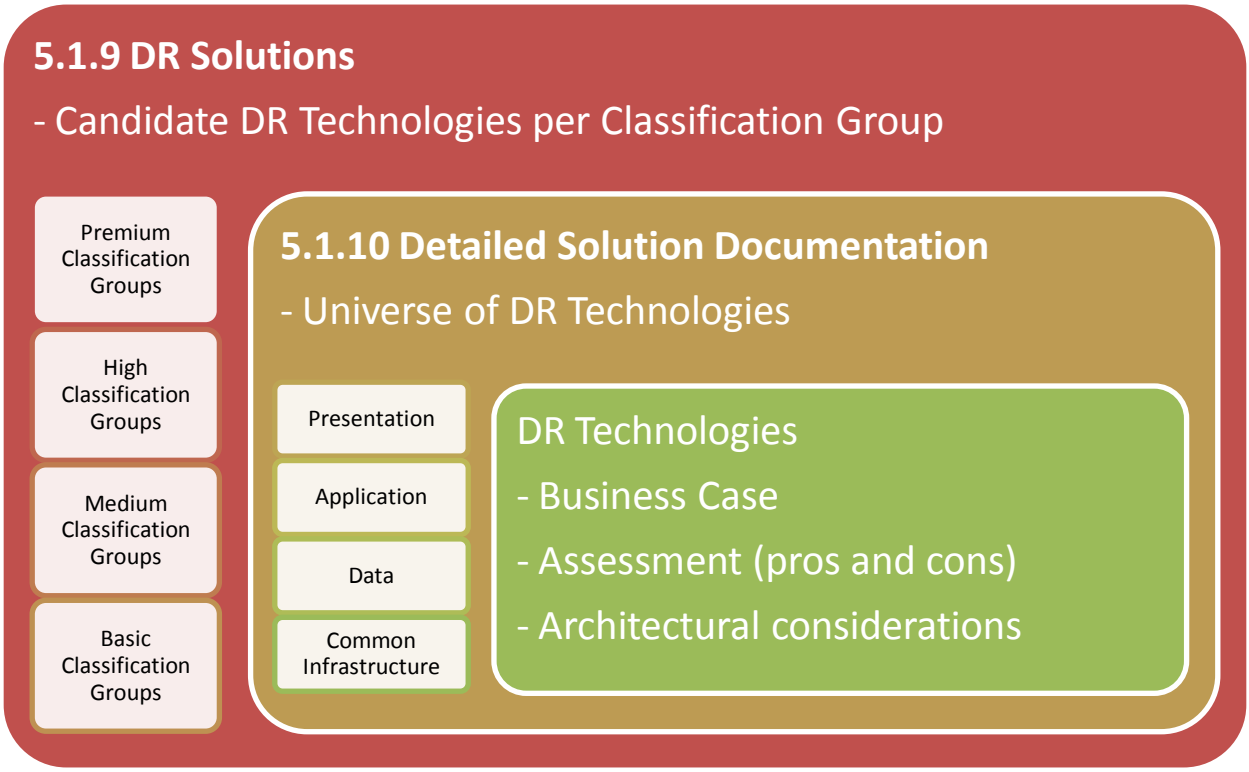

This document, 5.1.9, and 5.1.10 contain information relating to DR technology solutions. 5.1.9 contains the reasoning behind why certain technologies should be selected for a given set of requirements. It also shows the organization of technologies into solutions at the classification group level. 5.1.10 focuses on DR solution information at the technology component level. The innermost box in Figure 1 represents the detailed research that has been done on individual technologies. These detailed solutions are organized in a four-layered model and include a business case, assessment, considerations, and other factors that are pertinent when choosing a DR technology.

5 RECOMMENDED DR TECHNOLOGY SOLUTION

The following disaster recovery solutions provide a framework for designing standard solutions that meet the requirements of specific application groups. Figure 2 below shows the application classification groups by name[9]. Each column represents a service tier[10], for which the RTO and RPO are listed. Each row represents a platform, and the intersections form the classification groups. The requirements imposed by the service tier and platform drive the recommended DR solutions.

Figure 2. Classification Groups

	Premium (RTO- 15 min, RPO – 15 min)	High (RTO – 12 hrs, RPO – 2 hrs)	Medium (RTO – 48 hours, RPO – 24 hours)	Basic (RTO – 30 days, RPO – 7 days)
Cloud Eligible	**Premium-Cloud Eligible** X	**High-Cloud Eligible**	**Medium-Cloud Eligible**	**Basic-Cloud Eligible**
Common System	**Premium-Common System**	**High-Common System**	**Medium-Common System**	**Basic-Common System**
VistA	**Premium-VistA**	**High-VistA**	**Medium-VistA**	**Basic-VistA**
Mainframe	**Premium-Mainframe** X	**High-Mainframe**	**Medium-Mainframe**	**Basic-Mainframe**

X – DR solution is not recommended

There are two classification groups for which a DR solution is not recommended:

1) Premium-Cloud Eligible – A Premium Cloud-Eligible solution can be designed and delivered. However, to meet the Premium service requirements, the cloud DR infrastructure must be up and running continually. This is not a cost-effective design, and it runs against the principles that traditionally make cloud solutions so attractive.

[9] Further details can be found in reference "5.1.6 A&B: Classification Groups Research, Proposal, and Analysis"

[10] Further details can be found in references: VA SEDR12-0578, "5.1.5 B: Analysis Report Documenting Proposed Service Tiers"

2) Premium-Mainframe – A Premium-Mainframe solution is technically achievable. However, strict performance requirements must be met to implement it in a real world enterprise environment. Additionally, it does not support the strategic direction of VA and as such is not recommended.

VistA is recognized as a platform type given its importance and use within VA. There are two main components to VistA; the front-end (VistA FE) which provides the web interface and application layer, and the back-end (VistA BE) which hosts the database. From a DR solution perspective, the key feature of VistA is the use of the Caché database. The presence of a Caché database is significant since Caché data replication provides unique challenges. For this reason, the VistA back-end is the focus of any VistA DR solution discussion in this document.

Each DR solution is explored in more detail in the following sections. Figure 3 below shows a summary of the DR technology layer characteristics per classification group. A more detailed summary can be found in Appendix C.

Figure 3. DR Solutions Summary

	Premium Classification Groups	High Classification Groups	Medium Classification Groups	Basic Classification Groups
Presentation Layer	Servers processing data at both sites, global load balancing, and pre-configured DNS, firewall, and other components	Servers processing data at the primary site only	Servers processing data at the primary site only	DR site may not have the servers; infrastructure may have to be procured and configured prior to fail-over
Application Layer	Designed with fail-over in mind; may support data replication; servers processing data at both sites	Application may not be aware of DR; servers processing data at the primary site only	Servers processing data at the primary site only	DR site may not have the servers; infrastructure may have to be procured and configured prior to fail-over
Data Layer	Database and file system replication; only one site processing data Bi-directional replication only with data partition	Database and file system replication for VistA, cloud eligible and large databases; virtualized systems may use hypervisor replication mechanisms; SAN level replication for all other uses	Nightly disk based backups copied to the DR site using COTS tools; virtualized systems may use hypervisor integrated backups	Off-site weekly backups restored at the DR site
Common Infrastructure	Both sites are capable of processing data	Both sites are capable of processing data	DR site may have systems powered down	DR site may have systems powered down; DR site may not even exist in some cases

5.1 Premium Classification Group Solutions

5.1.1 Introduction

This section describes candidate disaster recovery technologies and considerations for the Premium classification groups. Considerations applicable to specific platforms will be called out as appropriate.

Two classification groups are addressed in this section:

- Premium-Common Systems
- Premium-VistA

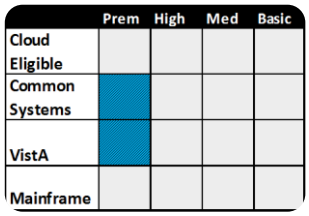

The Premium tier specifies a fifteen minute RTO and RPO. The common characteristics for this solution are the lack of manual intervention and the readiness of the DR infrastructure. The key features are described below:

- **Presentation Layer** − Servers processing workloads at primary and DR sites, automated failover with DNS-based redirection, global load balancing, and pre-configured user access components.
- **Application Layer** − Designed with failover to an active server in mind, may support data replication, and servers processing workloads at both sites.
- **Data Layer** – Asynchronous data replication. Options include:
    - File replication with managed file transfer technologies
    - Unidirectional replication at application or database level
    - Bidirectional replication only with data partitioning
    - VistA: Caché mirroring
- **Common Infrastructure** - Both primary and DR sites have computing, telecommunications, and environmental infrastructure needed to support services.

An active-active configuration is recommended for the Presentation and Application layers, meaning that backup servers are ready to, or already are, processing workloads. However, it is possible to have an active-passive configuration. In order to meet the RTO with an active-passive configuration, full automation of the failover process is recommended. Figure 4 below shows a view of the relevant disaster recovery technology components. The graphic represents each component at the layer in which it functions. Some icons represent specific technologies while others represent the technology features that are relevant in a DR scenario.

Figure 4. Premium Classification Group Technology Solutions

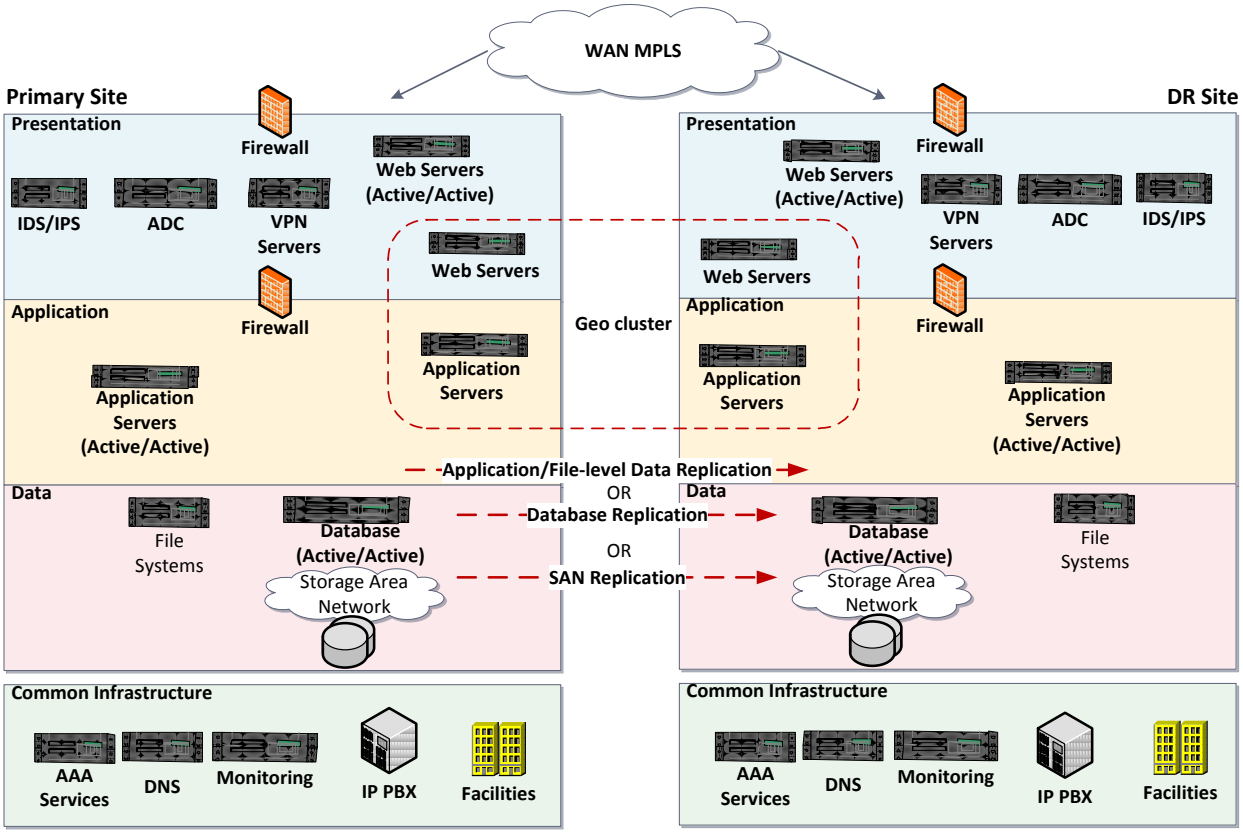

Although these components can support the RTO/RPO, they cannot guarantee the stated RTO and RPO unless the application itself has been architected with these goals in mind. The Premium solutions require redundant infrastructure and the ability to support near real time recovery. The cost to implement a Premium solution will almost certainly be expensive and should be justified through the BIA.

The ability to quickly recover from a disaster is a significant challenge that can require coordination of several layers of technology. Failover can be supported at all four layers of the solution, ranging from the use of active DR infrastructure, to virtual machine recovery software, to the redirection of workloads to the DR site. Automated failover refers to the use of IT tools to optimize the process of redirecting workloads to the DR infrastructure. However, there will be some additional cost, complexity, and risk involved in a fully automated solution. Examples of failover activities for live sites that will need to be automated are (Gregory 2007):

- Confirming the state of the data
- Making any required network configuration changes
- Promoting DR infrastructure to an active mode
- Testing DR application functionality
- Communicating the availability of recovered applications to users or consuming services

All such activities need to be automated, and failover may be orchestrated such that multiple systems are coordinated. The process of detecting a disaster should be carefully constructed to avoid the accidental triggering of an automated failover.

5.1.2 Presentation Layer

The presentation layer covers technologies that support client access to the services. Table 11 below presents technologies that are able to support the desired RTO and RPO, along with a recommendation for their inclusion (✓), exclusion (×), or consideration (?). References to more detailed information for each technology component can be found in Appendix C. At this layer, the solution components for the Premium-Common Systems and Premium-VistA groups are the same.

Table 11. Premium Classification Group – Presentation Layer Components

Component		Comments
Firewalls	✓	Necessary to meet security requirements.
IDS/IPS	✓	Necessary to meet security requirements.
VPN Servers	✓	Supports remote access by users.
Web Server Geo-Clustering	✓	Supports the RTO by not requiring a failover; the application design must support this use.
Active-Active Web Servers	✓	Supports the RTO by not requiring a failover; the application design must support this use.
HTTP Redirect	×	Can provide site selection function but is protocol specific and does not address all needs.
Layer 3 Route Health Injections	×	Can provide site selection function but does not provide intelligent load balancing capabilities. Routers and application aware devices must be integrated so health monitoring and route injection are implemented properly.
ADC (see below):	?	Provides various features including DNS-based redirection for access to services. Note that ADC is a term that includes WAN optimization, local load balancer, and GSLB.
• GSLB	✓	DNS-based traffic redirection functions that are provided by GSLBs produce better load sharing which is not available within DNS servers.
• Local Load Balancers	?	Should be used as needed to support services.
• WAN Optimization	?	Should be used as needed to support services.

Some key considerations are:

Active DR Infrastructure: An active presentation layer has been identified as a necessary feature for supporting the Premium service tier. The implications of this are as follows:

- DR Computing infrastructure such as web servers and user access components will be processing the workloads at any time. They should be treated as production equipment with regard to the ITSM systems and processes such as change management.
- Occasional failovers should verify the ability of DR infrastructure to function without the primary infrastructure supporting workload.

Session Persistence: An issue during failover is how to handle information that must be kept across the multiple requests in a user's session. In the event of a disaster, it is possible for this information to be

lost unless a mechanism for maintaining it is implemented. Possible solutions are the use of session databases, client-side cookies, or leveraging session persistence features of network elements.

5.1.3 Application Layer

The application layer covers technologies that are implemented at the application level. Table 12 below presents technologies that are able to support the desired RTO and RPO , along with a recommendation for their inclusion (✓), exclusion (✗), or consideration (?). References to more detailed information for each technology component can be found in Appendix C. These technologies are recommended for the Premium-Common Systems group only. No technologies are recommended at this layer for Premium-VistA.

Table 12. Premium Classification Group – Application Layer Components

Component		Comments
HA Application Clustering	✓	Use for Common Systems. Supports the RTO by not requiring a failover; the application design must support this use.
Active-Active Application Servers	✓	Use for Common Systems. Supports the RTO by not requiring a failover; the application design must support this use.

Some key considerations are:

Application Design: Active application servers at the DR site are necessary to support the RTO and are a key part of the DR solution. To leverage technologies such as active-active servers and application clustering, the application must be designed with these requirements in mind.

Internal and External Data Consistency: During the failback process, the state of application data must be reconciled between the original data sources and the disaster recovery sources, as well as with any consuming services.

5.1.4 Data Layer

The data layer covers technologies that support a solution's ability to meet the RPO by ensuring the availability of data. This is a critical portion of any disaster recovery solution. Table 13 below presents technologies that are able to support the desired RTO and RPO, along with a recommendation for their inclusion (✓), exclusion (✗), or consideration (?). References to more detailed information for each technology component can be found in Appendix C.

Table 13. Premium Classification Group – Data Layer Components

Component		Comments
Application-level Replication	✓	Use for Common Systems in place of or in conjunction with Database-level replication.
Managed File Transfer	✓	Supports replication of data stored in file systems.
Virtual Machine Replication	✗	Requires a restart of the system and may not meet RTO requirement.
Database-level Replication	✓	Supports the RTO/RPO.
Active-Active Database Configuration	✓	Use for Common Systems. Supports the RTO/RPO.
Storage-based Replication	✗	Imposes significant requirements to meet the RTO.

Storage area network (SAN) replication is a potential data solution for a fifteen minute RTO, but there are significant requirements with regard to computing and network infrastructure. Given the potential cost, it is not known how many applications will use a Premium tier solution. Database replication is recommended as it can be implemented on a per application basis. Managed file transfers should be used to replicate other forms of data.

It is vital to test the data replication solution to ensure that the data is valid and can be used in the alternate environment. One way that data testing can occur is during functional exercises, which are further described in Appendix B.

5.1.4.1 Classification Group Consideration: Common Systems

Asynchronous database replication is recommended for Common Systems. Two options may be used, based on the application design:

- **Active-Passive Databases with Database Replication:** One possible configuration is to have the active database at the primary location, where the writes are replicated to the DR database. The DR database is operating in a passive mode, where it is not used to process workloads until a disaster has occurred.
- **Active-Active Databases with Data Partitioning:** Database-level replication solutions sometimes offer features that enable the databases at remote locations to operate in an active-active configuration, where writes are possible at both ends. Examples include Oracle's Golden Gate solution and Microsoft's SQL Server replication with Merge Option. In these solutions, database log files shipped asynchronously to the DR sites are mined, and the database transactions are extracted and applied at the remote databases as if the transactions were executed on the remote database. These configurations permit near zero RTOs to be achieved, since there is no

requirement for a database failover in case of a disaster – the DR database is already running and processing data. RPO will be greater than zero, given that the replication is asynchronous. Data partitioning is implemented to minimize potential collisions between updates at both locations on the same data.

5.1.4.2 Classification Group Consideration: VistA

VistA data replication can be accomplished using features of the Intersystems Caché database. The Caché database is capable of asynchronous mirroring between two or more systems by pushing journal files, which reduces risks from out-of-order updates and carry-forward corruption.

Caché database mirroring can provide a disaster recovery solution using two different implementations. The first configuration allows for a special asynchronous member that can be configured to receive updates from multiple mirrored members, allowing for multiple mirrors to update one member. The second configuration allows for a mirror to provide constant updates to multiple asynchronous members so that one mirror can be backed up onto members in up to six different geographic locations.

The three key considerations for supporting this feature are:

- Network bandwidth and latency between the source and target locations must support the replication needs.
- The DR infrastructure must be sized appropriately to handle the workload.
- Redirection of traffic to the DR VistA instance must be handled by external means.

A newer release of Caché may be required to support these configurations; Caché release 2010 is recommended at a minimum.

5.1.5 Common Infrastructure Layer

The common infrastructure layer covers technologies that provide services across the computing infrastructure. While many of these technologies do not have specific disaster recovery functions, they are necessary to support Service Delivery and Service Assurance functions. Table 14 below presents technologies that are able to support the desired RTO and RPO, along with a recommendation for their inclusion (✓), exclusion (✗), or consideration (?). References to more detailed information for each technology component can be found in Appendix C.

Table 14. Premium Classification Group – Common Infrastructure Layer Components

Component		Comments
HA DNS	✓	Necessary to support Service Delivery/Assurance.
Directory Services	✓	Necessary to support Service Delivery/Assurance.
IP-PBX	✓	Classified as critical infrastructure by VA and supports telephone services.
Monitoring Systems	✓	Necessary to support Service Delivery/Assurance.
AAA Systems	✓	Necessary to support Service Delivery/Assurance.
LAN Technologies	✓	Supports connection via Local Area Network.
MPLS-WAN	✓	Supports connection via Wide Area Network.
Wireless Access Technologies	?	Use when conventional methods for provisioning network connections are difficult or do not meet time deadlines or other constraints.

Data Center Facilities	✓	Supports hosting of DR infrastructure.

Some key considerations are:

Live DR Facility: A live DR site is already supporting workloads and is ready to assume the workload of the primary site within the specified RTO. This implies the following:

- Network and environmental infrastructure, such as uninterruptible power supply (UPS) and HVAC, are sized to support additional workload.
- Computing infrastructure exists and is active or can be made active in time to support the additional workload.

Application Dependencies: Some common services, such as DNS or Directory Services, must be available before applications can begin the recovery process. These dependencies must be accounted for in the overall DR solution design to meet the stated RTO.

WAN Requirements: The network connection between the primary and DR site must be able to support the data replication/copy activities necessary to support a fifteen minute RPO. This may require dedicated, low-latency, high-bandwidth link or the use of optimization technologies.

5.2 High Classification Group Solutions

5.2.1 Introduction

This section describes candidate disaster recovery technologies and considerations for the High classification groups. Considerations applicable to specific platforms will be called out as appropriate.

Four classification groups are addressed in this section:

- High-Cloud Eligible
- High-Common Systems
- High-VistA
- High-Mainframe

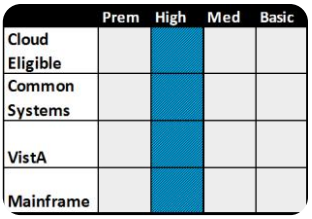

The High tier specifies a twelve hour RTO and a two hour RPO. The common characteristics for this solution are the semi-automated failover and the readiness of the DR infrastructure. The key features are described below:

- **Presentation Layer** − Servers and user access components are available and configured at the DR site, but not processing workloads. Failover is automated to the extent possible.
- **Application Layer** − Applications may not be aware of the DR solution; servers are available and configured at the DR site, but not processing workloads.
- **Data Layer –** Asynchronous data replication. Options include:
  - Common Systems: SAN replication, use of database replication if required; may use hypervisor replication
  - Mainframe: SAN replication
  - Cloud Eligible: Database replication, hypervisor replication, file transfers
  - VistA: Caché mirroring
- **Common Infrastructure** − Both primary and DR sites have computing, telecommunications, and environmental infrastructure needed to support services.

Figure 5 below shows a view of the relevant disaster recovery technology components.

Figure 5. High Classification Group Technology Solutions

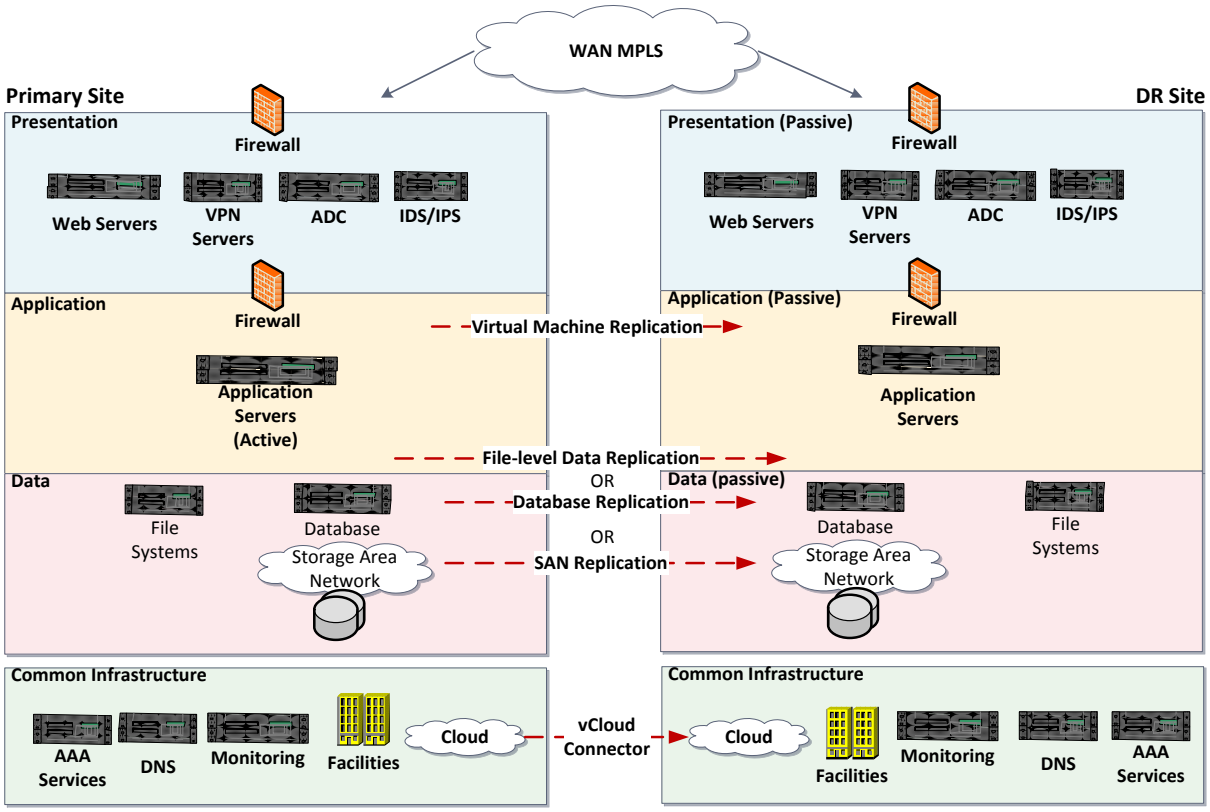

5.2.1.1 Solution-level Considerations

Although the DR infrastructure needs to be readily available to meet the recovery objectives, it is not required to be actively processing workloads. This means that computing infrastructure at the presentation, application, and data layers must be available and ready to be brought online at the DR site after declaration of disaster. To save operational costs, it is possible to power down some of the equipment. Some items, such as the SAN, may need to be kept on to participate in data replication activities. This passive state does not apply to the Common Infrastructure, which is described more fully later in this document.

The passive state of the DR infrastructure requires manual control over the failover. There will be manually executed processes but automation is still required to meet the twelve hour RTO. The use of IT tools to optimize the process of redirecting workloads to the DR infrastructure should be leveraged. Sample failover activities for this scenario are (Gregory 2007):

- Confirming the state of the data.
- Making any required network configuration changes.
- Promoting DR infrastructure to an active mode.

- Testing DR application functionality.
- Communicating the availability of recovered applications to users or consuming services.

Given the twelve hour RTO requirement, any configurations that require manual implementation should be in a ready state, meaning that the configuration has already been determined, tested, documented and is available. It also means that personnel must be available to make the needed configuration changes. The number of staff required to enact a DR solution must be accounted for as part of the application design requirements and implementation, as well as any other vendor or contractor support needed to meet the RTO.

The passive state of the DR computing infrastructure makes configuration management extremely important. Although the DR servers and components may not be processing workloads or even be powered on, it is still important they be kept up to date with production changes. These items and any related application updates must be included as part of the change/configuration management processes. The process for maintaining these servers is critical to ensuring that DR infrastructure remains in sync with the primary infrastructure.

5.2.1.1.1 Classification Group Consideration: High-Cloud Eligible

Applications that are determined to be Cloud Eligible will be using a public cloud solution. In most cases, the customer will not be able to specify the type of technology or infrastructure being used to provide the service. In this case, the DR infrastructure may not match the production infrastructure and may require different procedures to recover services. These potential differences must be well-understood and documented as part of the DR solution.

5.2.2 Presentation Layer

The presentation layer covers technologies that support client access to the services. Table 15 below presents technologies that are able to support the desired RTO and RPO, along with a recommendation for their inclusion (✓), exclusion (×), or consideration (?). References to more detailed information for each technology component can be found in Appendix C.

Table 15. High Classification Group – Presentation Layer Components

Component		Comments
Firewalls	✓	Necessary to meet security requirements.
IDS/IPS	✓	Necessary to meet security requirements.
VPN Servers	✓	Supports remote access by users.
Web Server Clustering	×	Imposes additional complexity and is not required to meet the RTO
Active-Active Web Servers	×	Imposes additional complexity and is not required to meet the RTO
Passive DR Web Servers	✓	DR servers must be available to meet the 12 hour RTO and are not processing workloads until disaster recovery.
HTTP Redirect	×	Can provide site selection function but is protocol specific and does not address all needs.
Layer 3 Route Health Injections	×	Can provide site selection function but does not provide intelligent load balancing capabilities, and routers and application aware devices must be integrated so health

Component		Comments
		monitoring and route injection are implemented properly.
ADC (see below)	?	Provides various features including DNS based redirection for access to services. Note that ADC is a term that includes WAN optimization, local load balancer, and GSLB.
• GSLB	?	Passive DR infrastructure does not require automated failover of client requests.
• Local Load Balancers	?	Use as needed within the application design.
• WAN Optimization	?	Should be used as needed to support services. Use for cloud eligible platforms should be considered.

A key consideration is the redirection of client access. At the High service tier, applications require some intervention before being able to provide service. Client access to these services should not be granted until they are available. A twelve hour RTO allows for a manual decision regarding when to allow client access to the DR infrastructure. However, proper preparation is needed to support this action. For example, a new firewall or load balancer configuration could be loaded to direct clients to the DR site, but this requires that the configuration be tested and available beforehand.

5.2.3 Application Layer

The application layer covers technologies that are implemented at the application level. Table 16 below presents technologies that are able to support the desired RTO and RPO, along with a recommendation for their inclusion (✓), exclusion (×), or consideration (?). References to more detailed information for each technology component can be found in Appendix C.

Table 16. High Classification Group – Application Layer Components

Component		Comments
Passive Application Servers	✓	DR servers must be available to meet the 12 hour RTO and are not processing workloads until disaster recovery.
HA Application Clustering	×	Imposes additional complexity and is not required to meet the RTO.
Active-Active Application Servers	×	Imposes additional complexity and is not required to meet the RTO.
vCloud Connector	?	Consider use for cloud eligible platforms which already exist in private or hybrid clouds.

To meet the twelve hour RTO, the application servers must be available and configured to process workloads.

5.2.4 Data Layer

The data layer covers technologies that support a solution's ability to meet the RPO by ensuring the availability of data. Table 17 below presents technologies that are able to support the desired RTO and RPO, along with a recommendation for their inclusion (✓), exclusion (×), or consideration (?). References to more detailed information for each technology component can be found in Appendix C.

Table 17. High Classification Group – Data Layer Components

Component		Comments
Application-level Replication	×	Storage-level replication provides robust enterprise level data replication architecture which obviates the need for application level replication in many situations.
Managed File Transfer	✓	Use for cloud eligible; supports replication of data stored in file systems.
Virtual Machine Replication	?	Consider use with cloud eligible and common system virtualized platforms.
Host-based Replication	×	Storage-level replication provides robust enterprise level data replication architecture which obviates the need for host based replication in many situations.
Database Replication	✓	Should be used for cloud eligible solutions as vendor cannot guarantee to match source SAN equipment, and for common systems with high volume databases.
Storage-based Replication	✓	Storage-level replication provides robust enterprise level data replication architecture which meets the RTO/RPO.
Mainframe Disaster Recovery	×	Mainframe specific solutions can add complexity; Storage-level replication provides robust enterprise level data replication architecture which obviates the need for mainframe DR solutions in many situations.

Each classification group may use a different data replication/copy solution as suited to the platform.

It is vital to test the data replication solution to ensure that the data is valid and can be used in the alternate environment. One way that data testing can occur is during functional exercises, which are further described in Appendix B.

5.2.4.1 Classification Group Consideration: High-Cloud Eligible

Data replication for Cloud Eligible platforms must be accomplished via database replication or host-based replication mechanisms that are independent of storage to de-couple the storage technology dependencies between the VA data center and the cloud vendor infrastructure.

Virtual machine configuration information should be copied to the cloud daily. A limited number of virtual machines will need to be constantly maintained on the cloud to enable this transfer. The typical bandwidth requirements for such copy operations are limited. State-aware cloud DR deployments require that business data be copied to the cloud. The two hour RPO would require that this data be transferred every two hours at a minimum.

5.2.4.2 Classification Group Consideration: High-VistA

For VistA systems, given the lack of clear integration between Caché databases and the underlying SAN storage, database replication provided by the Caché vendor is recommended. VistA data replication can be accomplished using features of the InterSystems Caché database or through SAN replication if the VistA instance is implemented with Linux. In the case of Linux, data replication should be performed via

the SAN as it would be for the mainframe and common systems platforms described in the next section. If the VistA instance is utilizing DEC Alpha hardware, Caché database replication is recommended.

VistA data replication can be accomplished utilizing features of the Intersystems Caché database. The Caché database is capable of asynchronous mirroring between two or more systems by pushing journal files, which reduces risks from to out-of-order updates and carry-forward corruption.

Caché database mirroring can provide a disaster recovery solution using two different implementations. The first configuration allows for a special asynchronous member that can be configured to receive updates from multiple mirrored members, allowing for multiple mirrors to update one member. The second configuration allows for a mirror to provide constant updates to multiple asynchronous members so that one mirror can be backed up onto members in up to six different geographic locations.

The three key considerations for supporting this feature are:

- Network bandwidth between the source and target locations must support the replication needs.
- The DR infrastructure must be sized appropriately to handle the workload.
- Redirection of traffic to the DR VistA instance must be handled by external means.

A newer release of Caché may be required to support these configurations; Caché release 2010 is recommended at a minimum.

5.2.4.3 Classification Group Consideration: High-Mainframe and High-Common Systems

SAN level replication is well suited for the High group as it is a common enterprise replication solution that achieves low RPOs and supports RTOs similar to the High service tier. It should be used where possible for the High tier.

In some cases, however, the application requirements make SAN replication an unsuitable solution. For example, for large databases with heavy input/output loads, SAN replication will require high levels of bandwidth. Database level replication based on log shipping can reduce the bandwidth requirements and should be considered as a secondary option in these situations. Database replication can be combined with SAN replication for file data replication.

5.2.5 Common Services Layer

The common services layer covers technologies that provide common services across the computing infrastructure. While many of these do not have specific disaster recovery functions, they are necessary to support Service Delivery and Service Assurance functions. Table 18 below presents technologies that are able to support the desired RTO and RPO, along with a recommendation for their inclusion (✓), exclusion (✗), or consideration (?). References to more detailed information for each technology component can be found in Appendix C.

Table 18. High Classification Group – Common Infrastructure Layer Components

Component		Comments
HA DNS	✓	Necessary to support Service Delivery/Assurance
Directory Services	✓	Necessary to support Service Delivery/Assurance
IP-PBX	×	Classified as critical infrastructure by VA and should be supported by a Premium solution
Monitoring Systems	✓	Necessary to support Service Delivery/Assurance
AAA Systems	✓	Necessary to support Service Delivery/Assurance
LAN Technologies	✓	Supports connection via Local Area Network
MPLS-WAN	✓	Supports connection via Wide Area Network
Wireless Access Technologies	?	Use when conventional methods for provisioning network connections are difficult or do not meet time deadlines or other constraints.
Data Center Facilities	✓	Supports physical hosting of DR infrastructure
Cloud	✓	Use for Cloud Eligible Platforms

Some key considerations are:

Network Capacity/Performance: When architecting the network for use in a DR scenario, one should keep in mind that there may be an increased load on the network above normal requirements. Additional data replication activities or user transactions need to be considered and accounted for.

Dependencies: Some common services, such as DNS or Directory Services, must be available before applications can begin the recovery process. Common infrastructure may be required to support applications in other service tiers. One likely scenario is that this infrastructure would be designed to meet the highest service tier being supported across the site and may already be included in higher level recovery plans. These dependencies must be accounted for in the overall DR solution design to meet the stated RTO.

5.3 Medium Classification Group Solutions

5.3.1 Introduction

This section describes candidate disaster recovery technologies and considerations for the Medium classification groups. Considerations applicable to specific platforms will be called out as appropriate.

Four classification groups are addressed in this section:

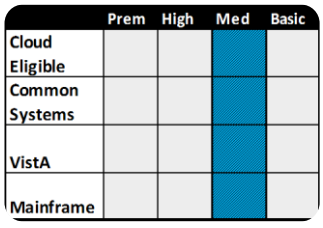

- Medium-Cloud Eligible
- Medium-Common Systems
- Medium-VistA
- Medium-Mainframe

The medium tier specifies a forty-eight hour RTO and twenty-four hour RPO. The key features of a disaster recovery solution that support this goal are:

- **Presentation Layer** – Servers are available and configured at the DR site, but may require manual configuration to process workloads. User access components may also require manual configuration.
- **Application Layer** – Servers are available and configured at the DR site, but may require manual configuration to process workloads.
- **Data Layer** – Nightly disk-based backups are copied to the DR site. Virtualized systems may use hypervisor integrated backups.
- **Common Infrastructure** – Computing, telecommunications, and environmental infrastructure exist at the DR facility. Additional provisioning may be required.

Each of these items will be discussed further in the context of the technology layer each occupies. Figure 6 below shows a view of the relevant disaster recovery technology components.

Figure 6. Medium Classification Group Technology Solutions

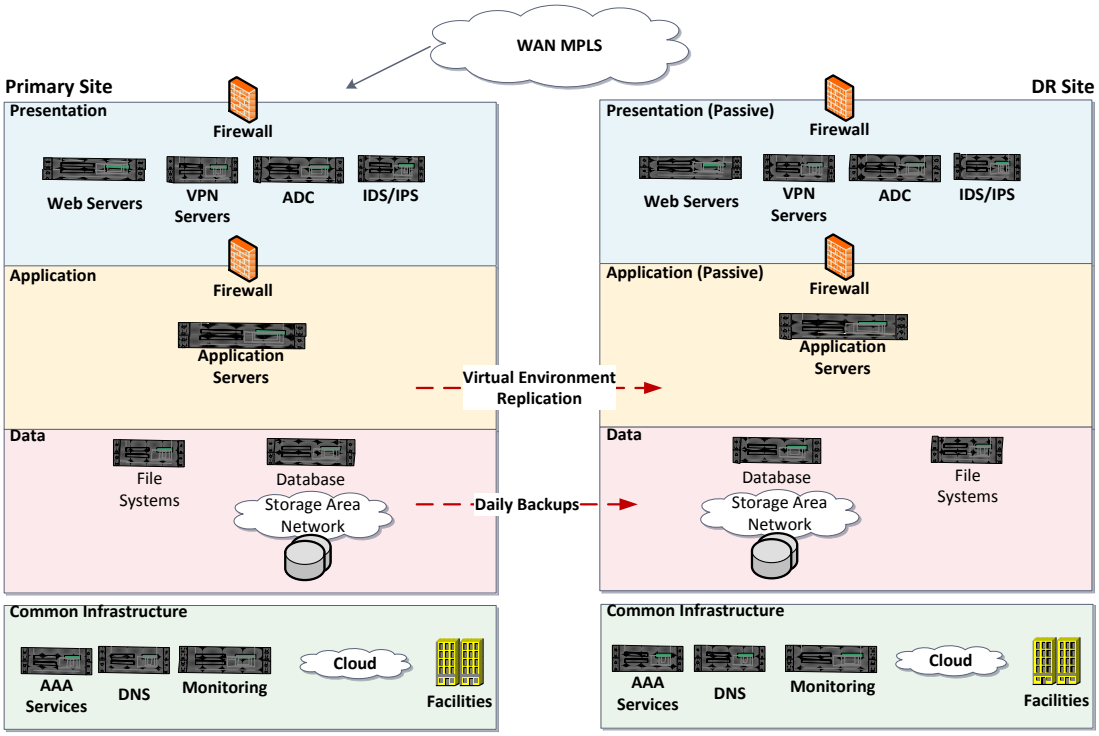

With a forty-eight hour RTO, manual configuration activities could include the complete building of servers, from installing the operating system, implementing patches, and installing and configuring applications, to restoring data. As with the High classification group solutions, the amount of staff and support needed to accomplish these activities must be accounted for in the design, especially for an enterprise data center that could host a large quantity of applications that need to be recovered. These activities also require that up-to-date versions of the software be available for use, possibly via a definitive media library, as well as access to the data backups for the systems. The availability of these repositories should also be accounted for in DR solution design.

The use of IT tools to optimize the process of redirecting workloads to the DR infrastructure should be leveraged. Sample failover activities associated with this scenario are (Gregory 2007):

- Retrieve the most recent backup media.
- Configure network devices and test connectivity.
- Install or update operating systems and applications on new servers.
- Restore data from backup media.
- Start the applications and perform functionality tests.
- Announce the availability of recovered applications as needed.
- Manually configure computing infrastructure

5.3.2 Presentation Layer

The presentation layer covers technologies that support client access to the services. Table 19 below presents technologies that are able to support the desired RTO and RPO, along with a recommendation for inclusion (✓), exclusion (✗), or consideration (?). References to more detailed information for each technology component can be found in Appendix C.

Table 19. Medium Classification Group – Presentation Layer Components

Component		Comments
Firewalls	✓	Necessary to meet security requirements.
IDS/IPS	✓	Necessary to meet security requirements.
VPN Servers	✓	Supports remote access by users.
Passive DR Web Servers	✓	DR servers must be available to meet the 48 hour RTO and are not processing workloads until disaster recovery.
HTTP Redirect	✗	Can provide site selection function but is protocol specific and does not address all needs.
Layer 3 Route Health Injections	✗	Can provide site selection function but does not provide intelligent load balancing capabilities, and routers and application aware devices must be integrated so health monitoring and route injection are implemented properly.
ADC	?	ADCs provide various features including DNS based redirection. Note that ADC is a term that includes WAN optimization, local load balancer, and GSLB. Use as needed to support failover and application performance.

While hardware must be available to support the RTO, the equipment may not be configured to support the workload from the primary site. It is strongly recommended that the hardware and software configurations be available and previously tested.

5.3.3 Application Layer

The application layer covers technologies that are implemented at the application level. Table 20 below presents technologies that are able to support the desired RTO and RPO, along with a recommendation for their inclusion (✓), exclusion (✗), or consideration (?). References to more detailed information for each technology component can be found in Appendix C.

Table 20. Medium Classification Group – Application Layer Technology Components

Component		Comments
Passive DR Application Servers	✓	DR servers must be available to meet the 48 hour RTO and are not processing workloads until disaster recovery.
HA Application Clustering	×	Imposes additional complexity and is not required to meet the RTO
Active-Active Application Servers	×	Imposes additional complexity and is not required to meet the RTO

Application servers must be available to support the RTO. However, they may require configuration before being able to process workloads. As with other equipment, it is strongly recommended that any software and configurations be ready and available.

5.3.4 Data Layer

The data layer covers technologies that support a solution's ability to meet the RPO. Table 21 below presents technologies that are able to support the desired RTO and RPO, along with a recommendation for their inclusion (✓), exclusion (×), or consideration (?). References to more detailed information for each technology component can be found in Appendix C.

Table 21. Medium Classification Group – Data Layer Components

Component		Comments
Virtual Machine Replication	?	Consider use with Cloud Eligible and Common System virtualized platforms.
Database-level Replication	×	RTO/RPO can be met with disk-based backups.
Storage-based Replication	×	RTO/RPO can be met with disk-based backups.
DR based on Backup (Disk/Tape)	✓	Disk-based backups recommended.

For Common Systems, VistA, and Mainframe, daily data backups are recommended. A possible exception exists for virtualized Common Systems and the Cloud Eligible platform, as detailed in the next sections.

It is vital to test the data replication solution to ensure that the data is valid and can be used in the alternate environment. One way that data testing can occur is during functional exercises, which are further described in Appendix B.

5.3.4.1 Classification Group Considerations: Medium-Common Systems

For virtualized Common Systems, virtual environment replication can be used to perform data replication/copy functions. The virtual machine information can also be used to recreate web and application servers at the DR site.

5.3.4.2 Classification Group Considerations: Medium-Cloud Eligible

Virtual machine configuration information should be copied to the cloud daily. A limited number of virtual machines are constantly maintained on the cloud to enable this transfer. Typically the bandwidth

requirement for such copy operations is limited. State-aware cloud DR deployments require that business data also be copied to the cloud. For the Medium service tier, data may be copied over nightly to achieve the twenty-four hour RPO.

5.3.5 Common Infrastructure Layer

The common services layer covers technologies that provide common services across the computing infrastructure. Although many of these do not have specific disaster recovery functions, they are necessary to support Service Delivery and Service Assurance functions. Table 22 below presents technologies that are able to support the desired RTO and RPO, along with a recommendation for their inclusion (✓), exclusion (✗), or consideration (?). References to more detailed information for each technology component can be found in Appendix C.

Table 22. Medium Classification Group – Common Infrastructure Layer Components

Component		Comments
HA DNS	✓	Necessary to support Service Delivery/Assurance
Directory Services	✓	Necessary to support Service Delivery/Assurance
IP-PBX	✗	Classified as critical infrastructure by VA and should be supported by a Premium solution
Monitoring Systems	✓	Necessary to support Service Delivery/Assurance
AAA Systems	✓	Necessary to support Service Delivery/Assurance
LAN Technologies	✓	Supports connection via Local Area Network
MPLS-WAN	✓	Supports connection via Wide Area Network
Wireless Access Technologies	?	Use when conventional methods for provisioning network connections are difficult or do not meet time deadlines or other constraints.
Data Center Facilities	✓	Supports physical hosting of DR infrastructure
Cloud	✓	Use for Cloud Eligible Platforms

Similar to the High classification group solution, some common services, such as DNS or Directory Services, must be available before applications can begin the recovery process. Common infrastructure may be required to support applications in other service tiers. One likely scenario is that this infrastructure would be designed to meet the highest service tier being supported across the site and may already be included in higher level recovery plans. These dependencies must be accounted for in the overall DR solution design to meet the stated RTO.

5.4 Basic Classification Group Solutions

5.4.1 Introduction

This section describes candidate disaster recovery technologies and considerations for the Basic classification groups. Considerations applicable to specific platforms will be called out as appropriate.

Four classification groups are addressed in this section:

- Basic-Cloud Eligible
- Basic-Common Systems
- Basic-VistA

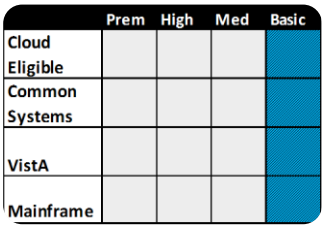

- Basic-Mainframe

The Basic tier specifies a thirty day RTO and seven day RPO. At this level, much of the DR infrastructure may not exist and will require procurement activities. As such, it is usually the least costly type of DR solution because there is very little upfront investment. The key features of a disaster recovery solution that support this goal are:

- **Presentation Layer** – Infrastructure may not exist and will have to be procured. Servers and user access components will require configuration to process workloads.
- **Application Layer** – Servers may not exist and will have to be procured. Servers will require configuration to process workloads.
- **Data Layer** – Off-site weekly backups.
- **Common Infrastructure** – DR facility, environmental, and network infrastructure may not exist and will have to be procured. Provisioning and configuration activities may be required.

Figure 7 below shows a view of the relevant disaster recovery technology components.

Figure 7. Basic Classification Group Technology Solutions

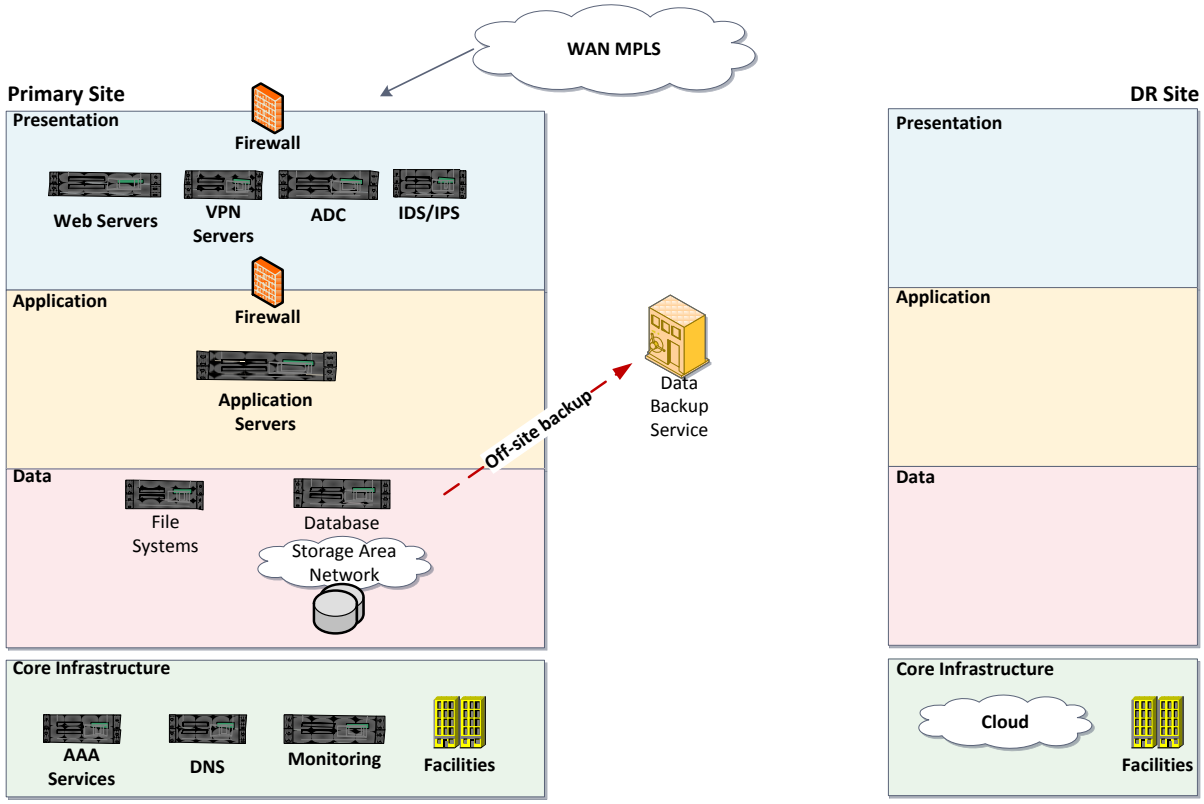

The common theme of the Basic classification group solutions is the lack of DR infrastructure, and manual failover. Computing infrastructure such as servers and networking equipment need to be procured. To meet the thirty day RTO, pre-existing agreements with vendors are recommended to

expedite the procurement process. Drop shipping services may be another way to expedite the procurement process.

Sample failover activities for this scenario are (Gregory 2007):

- Order systems and network devices to replace the systems and devices damaged by the disaster.
- Retrieve the most recent backup media.
- Install network devices and construct the network. Verify connectivity.
- Install operating systems and applications on new servers.
- Restore data from backup media.
- Start the applications and perform functionality tests.
- Announce the availability of recovered applications as needed.

5.4.2 Presentation Layer

The presentation layer covers technologies that support client access to the services. The components at this level would need to be ordered and set up to restore services. No DR-specific technologies are recommended.

5.4.3 Application Layer

The application layer covers technologies that are implemented at the application level. The components at this level would need to be ordered and set up to restore services. No DR-specific technologies are recommended.

5.4.4 Data Layer

The data layer covers technologies that support a solution's ability to meet the RPO. Table 23 below presents technologies that are able to support the desired RTO and RPO, along with a recommendation for their inclusion (✓), exclusion (✗), or consideration (?). References to more detailed information for each technology component can be found in Appendix C.

Table 23. Basic Classification Group – Data Layer Components

Component		Comments
DR based on Backup (Disk/Tape)	✓	Off-site backups are recommended to meet the RTO/RPO.

Off-site backups should be performed weekly and transferred off-site. Backups may be implemented via magnetic tape, optical storage, or electronic vaulting. Regular testing of the back up media should be performed to verify their use for recovery.

5.4.4.1 Basic-Cloud Eligible

The Basic-Cloud Eligible solution may consider backing up to a cloud provider. Note that no virtual machines are purchased as part of this backup. Costs will be incurred solely through use of the cloud vendor's data transfer and data storage services.

5.4.5 Common Infrastructure Layer

The common infrastructure layer covers technologies that provide common services across the computing infrastructure. Although many of these do not have specific disaster recovery functions, they are necessary to support Service Delivery and Service Assurance functions. Table 24 below presents technologies that are able to support the desired RTO and RPO, along with a recommendation for their inclusion (✓), exclusion (✗), or consideration (?). References to more detailed information for each technology component can be found in Appendix C.

Table 24. Basic Classification Group – Common Services Components

Common Services Component		Comments
Data Center Facilities	✓	Necessary to support hosting of DR infrastructure
Cloud	✓	Use for Cloud Eligible Platforms

An empty DR site is relatively inexpensive to operate as it does not contain any computing infrastructure. However, it does require that all infrastructure and telecommunications be provisioned at the time of disaster. Another possibility is not having a DR facility at all, which would require that a location also be procured at the time of disaster. Locating a facility with the required infrastructure, power, communications, and security is an activity that should be performed in advance. Because set up time may be lengthy, pre-existing arrangements need to be made with vendors, utility and service providers. Considerations for this scenario are:

- **Mobile computing facilities:** Mobile data centers are all-in-one container solutions that can be ordered with the required infrastructure and hardware. They can be installed anywhere there is an available power source and communication lines. Additional configuration is required after the mobile data center is in place.
- **Drop Shipping:** Drop shipping refers to services that specialize in the shipment of computing infrastructure directly to the customer. It is one method of acquiring equipment.
- **Hosted or Colocation Facilities:** The DR facility can be supplied by a third party. Hosting providers offer space, environmental, and network infrastructure to multiple customers. Security services are also available.

Areas that need to be addressed when procuring or establishing a new processing facility are:

- **Geographic location:** Avoiding areas that are known to be prone to natural disasters can limit cost and avoid potential loss of the DR facility due to a natural disaster.
- **Physical Security:** It is important to ensure the physical security of the data center infrastructure and to prevent people from trespassing.
- **Power Supply:** Commercial power, UPS, or generators.
- **Environmental Controls:** The environmental controls within the data center must be capable of maintaining the required temperature and humidity. When hardware is operated outside of the recommended environmental guidelines, there is a risk of failure that may lead to an outage.
- **Racks:** Racks maximize the quantity and volume of hardware that can be deployed in a prescribed amount of space and allow for the efficient organization of the hardware. When installed correctly, racks also provide the necessary electrical grounding.

- **Communication Lines:** Communication lines are required to get information both in and out of the data center. The standard lead times for both voice and data line orders normally exceed 30 days, which is why in the event of a disaster it is critical to already have a plan for both voice and data communications.

6 CONCLUSION

To support VA's needs in establishing enterprise-wide standards for disaster recovery, research and analysis of industry best practices was performed. This included examining VA and federal standards, examining input from VA stakeholders, and leveraging internal resources with relevant experience. This document has recommended four standard platforms that can be used along with service tiers to classify applications into groups. These groups can then be matched to the DR solutions recommended in this document, which provide a technology selection roadmap. This methodology should result in a standard set of enterprise-level DR solutions, which supports VA's goals of optimizing resources and increasing interoperability.

Appendix A. Security Controls

Several federal standards and policies are relevant to VA in the context of this project: FIPS 199, FIPS 200, NIST SP 800-34, and NIST SP 800-53. Additionally, VA Handbooks 6500, 6500.5, and 6500.8 address how these 800-53 controls are applied and implemented for systems and applications developed for, or used by, VA. These controls are especially pertinent during design and implementation phases, and for this reason they are further described here. The additional information is provided to assist the reader in understanding and accounting for future activities that will be required in solution implementation.

FIPS 199 and FIPS 200 deal specifically with categorization of systems (low, moderate, high). NIST SP 800-34 addresses Information Systems Contingency Planning (ISCP) and draws its requirements directly from NIST SP 800-53, which specifies criteria required for security control classes, families, and identifiers by class (i.e., Managerial, Operational, or Technical). Table A-1 below lists these security control items.

Table A-1. Security Control Classes, Families, and Identifiers

Identifier	Family	Class
AC	Access Control	Technical
AT	Awareness and Training	Operational
AU	Audit and Accountability	Technical
CA	Security Assessment and Authorization	Management
CM	Configuration Management	Operational
CP	Contingency Planning	Operational
IA	Identification and Authentication	Technical
IR	Incident Response	Operational
MA	Maintenance	Operational
MP	Media Protection	Operational
PE	Physical and Environmental Protection	Operational
PL	Planning	Management
PS	Personnel Security	Operational
RA	Risk Assessment	Management
SA	System and Services Acquisition	Management
SC	System and Communications Protection	Technical
SI	System and Information Integrity	Operational
PM	Program Management	Management

Specific control family identifiers prescribed in VA 6500 are: SA-3 Life Cycle Support, PL-1 Security Planning Policies and Procedures, SA-8 Engineering Principles as documented in NIST SP 800-27, SA-3

Information System Connections, PE-16 Delivery and Removal, CM-8 Information System Component Inventory, CA-3, PL-5 Privacy Impact, PL-2 SSP, RA-5 Vulnerability Scanning, and SI-2 Flaw Remediation.

The Federal Enterprise Architecture Security and Privacy Profile also prescribes specific control family identifiers. They are: CM-6 Configuration Settings, CM-2 Baseline configuration, VA Handbook 6500.3 (Certification and Accreditation of VA information systems), CA-2 Security Assessments, CA-5 Plans of Action and Milestones, CA-6 Security Authorization, CA-7 Continuous Monitoring, SI-4 Information System Monitoring Tools and Techniques, RA-5 Vulnerability Scanning, AU-6 Audit Monitoring, Analysis, and Reporting, SI-7 Software and Information Integrity, IR-4 Incident Response, CM-4 Monitoring Configuration Changes, MP-5 Media Transport, MP-6, Media Sanitization and Disposal, and CM-8 Information System Component Inventory.

Appendix B. Disaster Recovery Testing

TEST, TRAINING, AND EXERCISE (TT&E) PROGRAM

An effective TT&E program is necessary to prepare and validate VA's capabilities to continue to perform Primary Mission Essential Functions (PMEFs) and Mission Essential Functions (MEFs) during a significant disruptive event that results in the loss of key personnel, denial of facility access, or an IT/infrastructure outage.

Federal Continuity Directive (FCD) 1 describes those specific activities and phases needed to ensure the continuity of headquarters functions and their supporting infrastructure, while NIST SP 800-34 (which parallels NIST SP 800-53) addresses Information Systems Contingency Plan (ISCP) capabilities for all General Support Systems (GSS) and Major Applications (MAs). These are summarized below.

The testing, training, and exercising of continuity capabilities is essential to demonstrating, assessing, and improving VA's ability to execute its continuity program, plans, and procedures. **Training** familiarizes continuity personnel with their roles and responsibilities in support of the performance of an agency's essential functions during a continuity event. **Tests and exercises** serve to assess, validate, or identify for subsequent correction, all components of continuity plans, policies, procedures, systems, and facilities used in response to a continuity event. Periodic testing also ensures that equipment and procedures are kept in a constant state of readiness. The TT&E program should be part of a multiyear TT&E plan that addresses continuity TT&E requirements, resources to support TT&E activities, and a TT&E planning calendar. The following details the specific requirements for each component.

Testing and Exercising

The following testing and exercising areas should be addressed:

Annually:

- Alert, notification, and activation procedures for continuity personnel (Note: quarterly for HQ personnel)
- Plans for recovering vital records, critical information systems, services, and data.
- Primary and backup infrastructure systems and services (e.g., power, water, fuel) at alternate facilities.
- Required physical security capabilities at alternate facilities.
- Capabilities required to perform all MEFs, as identified in the business process analysis.
- Internal and external interdependencies identified in the plan, with respect to performance of all MEFs.
- Validation of the process for formally documenting and reporting tests and their results as directed.

Monthly:

- Testing and validating of equipment to ensure the internal and external interoperability and viability of communications systems, through monthly testing of the continuity communications.

ISCP testing enables plan deficiencies to be identified and addressed by validating one or more of the system components and the operability of the plan. Testing can take several forms and accomplish several objectives but should be conducted in as close to an operating environment as possible. Each information system component should be tested to confirm the accuracy of individual recovery procedures. The following areas should be addressed in a contingency plan test, as applicable:

- Notification procedures
- System recovery on an alternate platform from backup media
- Internal and external connectivity and dependencies
- System performance using alternate equipment
- Restoration of normal operations; and
- Other plan testing (where coordination is identified, e.g., COOP, Business Continuity Plan (BCP).

To derive the most value from the test, the ISCP Coordinator should develop a test plan designed to examine the selected element(s) against explicit test objectives and success criteria. The use of test objectives and success criteria enables the effectiveness of each system element and the overall plan to be assessed. The test plan should include a schedule detailing the time frames for each test and test participants. The test plan should also clearly delineate scope, scenario, and logistics. The scenario chosen may be a worst-case incident or an incident most likely to occur. It should mimic reality as closely as possible.

Tests are evaluation tools that use quantifiable metrics to validate the operability of an information system or system component in an operational environment. For example, an organization could test call tree lists to determine if calling can be executed within prescribed time limits; another test might be to remove power from a system or system component. A test is conducted in as close to an operational environment as possible; if feasible, an actual test of the components or systems used to conduct daily operations for the organization should be used. The scope of testing can range from individual system components or systems to comprehensive tests of all systems and components that support an ISCP. Tests often focus on recovery and backup operations; however, testing varies depending on the FIPS 199 impact level, the goal of the test, and its relation to a specific ISCP.

It is vital to test data backups to ensure they are valid and can actually work in the alternate environment. Testing can be performed in a modular fashion. For example, equipment and backups can be tested separately then together. However, it is essential that all components are ultimately tested independently and combined. One place where data testing can occur is during functional exercises for moderate and high-impact systems, which are described next.

NIST SP 800-84 identifies the following types of exercises widely used in information system TT&E programs by single organizations:

- **Tabletop Exercises –** Discussion-based only, and do not involve deploying equipment or other resources. A facilitator presents a scenario and asks the exercise participants

questions related to the scenario, which initiates a discussion among the participants of roles, responsibilities, coordination, and decision making.

- **Functional Exercises –** Allow personnel to validate their operational readiness for emergencies by performing their duties in a simulated operational environment. They:
  - Are designed to exercise the roles and responsibilities of specific team members, procedures, and assets involved in one or more functional aspects of a plan (e.g., communications, emergency notifications, system equipment setup)
  - Vary in complexity and scope, from validating specific aspects of a plan to full-scale exercises that address all plan elements
  - Allow staff to execute their roles and responsibilities as they would in an actual emergency situation, but in a simulated manner.

For low-impact systems, a tabletop exercise at an organization-defined frequency is sufficient. The tabletop should simulate a disruption, include all main ISCP points of contact, and be conducted by the system owner or responsible authority.

For moderate-impact systems, a functional exercise at an organization-defined frequency should be conducted. The functional exercise should include all ISCP points of contact and be facilitated by the system owner or responsible authority. Exercise procedures should be developed to include an element of system recovery from backup media.

For high-impact systems, a full-scale functional exercise at an organization-defined frequency should be conducted. The full-scale functional exercise should include a system failover to the alternate location. This could include additional activities such as full notification and response of key personnel to the recovery location, recovery of a server or database from backup media or setup, and processing from a server at an alternate location. The test should also include a full recovery and reconstitution of the information system to a known state.

Table B-1 below presents a sample of TT&E activity using NIST Special Publication 800-53 guidance and as required by the FIPS 199 impact level.

Table B-1: ISCP TT&E Activities

TT&E Event	Sample Activity	FIPS 199 Availability Security Objective
ISCP Training (CP-3)	A seminar and/or briefing used to familiarize personnel with the overall ISCP purpose, phases, activities, and roles and responsibilities.	Low Impact = Yes Mod. Impact = Yes High Impact = Yes

Instruction (CP-3)	Instruction of contingency personnel on their roles and responsibilities within the ISCP. Includes refresher training. (For a high-impact system, incorporate simulated events.)	Low Impact = Yes Mod. Impact = Yes High Impact = Yes
Contingency Plan Test/Exercise (CP-4)	Test and/or exercise the contingency plan to determine effectiveness and the organization's readiness. This could include planned and unplanned maintenance activities	All

Training

The training program must include (all performed at least annually):

- Continuity awareness briefings (or other means of orientation) for the entire workforce.
- Personnel (including host or contractor personnel) who are assigned to activate, support, and sustain continuity operations.
- Leadership on that agency's PMEFs and MEFs, including training on their continuity responsibilities.
- Agency personnel who assume the authority and responsibility of leadership if that leadership is incapacitated or becomes otherwise unavailable during a continuity situation.
- Pre-delegated authorities for making policy determinations and other decisions, at the headquarters, field, satellite, and other organizational levels, as appropriate.
- Personnel briefings on agency continuity plans that involve using, or relocating to, alternate facilities, existing facilities, or virtual offices.
- Capabilities of communications and IT systems to be used during a continuity event.
- Identification, protection, and ready availability of electronic and hardcopy documents, references, records, information systems, and data management software and equipment (including classified and other sensitive data) needed to support essential functions during a continuity situation.
- Devolution option for continuity, to address how each will identify and conduct its essential functions during an increased threat situation or in the aftermath of a catastrophic emergency.
- All reconstitution plans and procedures to resume normal operations from the original or replacement primary operating facility.

Appendix C. Disaster Recovery Solutions Summary

Table C-1 below summarizes the DR solution technology solution components by classification group. The DR technologies are listed down the left side and organized by layer: Presentation, Application, Data, and Common Services. Each column is labeled with the name of a classification group and contains an indication of the DR technology use. Where available, a reference to the detailed solution information provided by reference document "5.1.10 B: Detailed Solution Documentation. The table legend is as follows:

✓ This technology is recommended

? This technology should be considered for use based on application requirements

Table C-1. Classification Group DR Solution Components

	5.1.10 Section	Premium-Common Systems	Premium-VistA	High-Cloud Eligible	High-Common Systems	High-VistA	High Mainframe	Medium-Cloud Eligible	Medium-Common Systems	Medium-VistA	Medium-Mainframe	Basic-Cloud Eligible	Basic-Common Systems	Basic-VistA	Basic-Mainframe
Firewalls	2.1.1	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓				
IDS/IPS	2.1.2	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓				
VPN	2.1.3	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓				
Web Server Geo-cluster	2.2.1	✓	✓												
Active-Active Web Server	2.2.2	✓	✓												
Active-Passive Web Server				✓	✓	✓	✓	✓	✓	✓	✓				
ADC	2.3.1	?	?	?	?	?	?	?	?	?	?				
GSLB	2.3.2	✓	✓	?	?	?	?	?	?	?	?				
Load Balancers	2.3.3	?	?	?	?	?	?	?	?	?	?				
WAN Optimization	2.3.4	?	?	?	?	?	?	?	?	?	?				
HA Application Clustering	3.1	✓													
Active-Active Application Servers	3.2	✓													
Active-Passive Application Servers				✓	✓	✓	✓	✓	✓	✓	✓				
vCloud Connector	3.3			✓											
Application-level Replication	4.1.1	✓													
Managed File Transfer	4.1.2	✓	✓	✓		?									
Virtual Machine Replication	4.1.3			✓	?			?	?						
Host-based Replication	4.2.1														
Database Level Replication	4.3.1	✓	✓	✓	✓										
Active-Active Database Configurations	4.3.2	✓													
Storage-based Replication	4.4.1			✓	✓	✓	✓								
DR based on Backup (Disk/Tape)	4.4.2							✓	✓	✓	✓	✓	✓	✓	✓
Mainframe DR	4.4.3														
HA DNS	5.1.1	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓				
Directory Services	5.1.2	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓				
IP-PBX	5.1.3	✓	✓												
Monitoring Systems	5.1.4	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓				
AAA Services	5.1.5	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓				
LAN	5.2.1	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓				
MPLS-WAN	5.2.2	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓				
Wireless Access Technologies	5.2.3	?	?	?	?	?	?	?	?	?	?				
Data Center Facilities	5.3	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Cloud	5.4			✓				✓				✓			

List of Acronyms

Term	Definition
AAA	Authentication, Authorization, and Accounting
ADC	Application Delivery Controller
BIA	Business Impact Analysis
BCP	Business Continuity Plan
CICS	Customer Information Control System
COOP	Continuity of Operations
COTS	Commercial Off-the-Shelf
CPU	Central Processing Unit
DNS	Domain Name Service
DR	Disaster Recovery
ESE	Enterprise Systems Engineering
FCD	Federal Continuity Directive
FDCCI	Federal Data Center Consolidation Initiative
FIPS	Federal Information Processing Standard
GSLB	Global Server Load Balancing
GSS	General Support Systems
HA	High Availability
HTTP	Hypertext Transfer Protocol
HVAC	Heating, Ventilation, Air Conditioning
IDMS	Integrated Database Management System
IDS/IPS	Intrusion Detection System/Intrusion Prevention System
IP PBX	Internet Protocol Private Branch Exchange
ISCP	Information Systems Contingency Planning
IT	Information Technology
LAN	Local Area Network

MA	Major Application
MEF	Mission Essential Functions
MPLS	Multiprotocol Label Switching
NIST	National Institute of Standards and Technology
OIT	Office of Information and Technology
OS	Operating System
PMEF	Primary Mission Essential Functions
RPO	Recovery Point Objective
RTO	Recovery Time Objective
SAN	Storage Area Network
TRM	Technical Reference Model
TT&E	Test, Training, and Exercise
UPS	Uninterruptible Power Supply
U.S.	United States
VA	Department of Veterans Affairs
VistA	Veterans Health Information Systems and Technology Infrastructure
VistA FE	Veterans Health Information Systems and Technology Infrastructure Front End
VistA BE	Veterans Health Information Systems and Technology Infrastructure Back End
VPN	Virtual Private Network
WAN	Wide Area Network
AITC	Austin Information Technology Center
BCM	Business continuity management
CDCO	VA Corporate Data Center Operations
MTD	Maximum tolerable downtime
SEDR	Systems Engineering and Design Review
SMART	Security Management and Reporting Tool database

VBA	Veterans Benefits Administration
VHA	Veterans Health Administration

References

Baker, Mark, Charles Witschorik, Jonathan Tuch, Waverly Hagey-Espie, and Veena Mendiratta. "Architectures and Disaster Recovery Strategies for Survivable Telecommunications Services." *Bell Labs Technical Journal.* 2004. http://onlinelibrary.wiley.com/.

Cisco Systems. "Data Center Disaster Recovery and Business Continuance." 2009. http://www.ciscolive.com/global/virtual/ (accessed April 2, 2012).

Cisco Systems. *Disaster Recovery: Best Practices.* White Paper, Cisco, 2008.

Cohen, Lloyd, and Kuba Stolarski. *Wordwide and U.S. Enterprise Server 2010 Vendor Shares: x86 Shipments and Installed Base.* Competitive Analysis, IDC, 2011.

Corporate Data Center Operations. *2010 Brochure.* 2010. http://www.cdco.va.gov/docs/2010_Brochure.pdf (accessed 05 17, 2012).

Defense Information Systems Agency, Computing Services Directorate. *Catalog of Services.* May 4, 2011.

Department of Veterans Affairs. "Draft ESLA 2.11 ISCP and DR 17 Feb 2012." February 2012.

Department of Veterans Affairs, Office of Information and Technology, EIE. *VA OI&T Virtualization Plaftorms Procurement Guidelines.* March 22, 2010.

Department of Veterans Affairs, Office of Information and Technology, Enterprise Infrastructure Engineering. *VA Enterprise IT Infrastructure Standard Server Platform Production V1.0.* November 18, 2009.

Department of Veterans Affairs, Office of Information and Technology, Service Delivery and Engineering (SDE). "VA Enterprise Disaster Recovery Service Tiers Standard, Version 1.0." VA Disaster Recovery Analysis, Task Order: VA118A-11-0178, SEDR12-0578, 2012.

Department of Veterans Affairs, Office of Information and Technology, Service Delivery and Engineering (SDE). "VA Enterprise Disaster Recovery Technology Solutions Standard, Version 1.0." VA Disaster Recovery Analysis, Task Order: VA118A-11-0178, 2012.

Department of Veterans Affairs, OIT Service Delivery and Engineering. *OIT Release Architecture V1.21.* November 30, 2011.

Google. *Disaster Recovery by Google.* 2010. http://googleenterprise.blogspot.com/2010/03/disaster-recovery-by-google.html .

Gregory, Peter. *IT Disaster Recovery Planning for Dummies.* Wiley Publishing, Inc., 2007.

InterSystems. "InterSystems Cache Database Mirroring." *InterSystems Cache Database Mirroring: Technical Summary.* n.d. http://www.intersystems.com/highavailability/Cache_Database_Mirroring_Technical_Overview.pdf (accessed June 7, 2012).

Mell, Peter, and Timothy Grance. *The NIST Definition of Cloud Computing.* NIST Special Publication 800-145, National Institute of Standards and Technology, 2011.

Needleman, Rafe. *CNET.* September 7, 2011. http://news.cnet.com/8301-17939_109-20102953-2/google-docs-suffers-30-minute-outage/ (accessed April 1, 2011).

Ray, Ashish, Mano Malayanur, and Dingbo Zhou. "The Right Choice for DR: Data Guard, Stretch Clusters, or Remote Mirroring." *Oracle Open World.* FannieMae, 2009.

Scaramella, Jed, et al. *Worldwide and Regional Server 2011-2015 Forecast.* Market Analysis, IDC, 2011.

Swanson, M, P Bowen, W Amy, D Gallup, and D Lynes. *NIST Special Publication 800-34, Rev. 1, Contingency Planning Guide for Federal Information Systems.* NIST, 2010.

The MITRE Corporation. "5.1.10 A: Detailed Solution Documentation Formats." VA Disaster Recovery Analysis, Task Order: VA118A-11-0178, 2012.

The MITRE Corporation. "5.1.10 B: Detailed Solution Documentation." VA Disaster Recovery Analysis, Task Order: VA118A-11-0178, 2012.

The MITRE Corporation. "5.1.3: Inventory and Discovery." VA Disaster Recovery Analysis, Task Order:VA118A-11-0178, 2012.

The MITRE Corporation. "5.1.5 B: Analysis Report Documenting Proposed Service Tiers." VA Disaster Recovery Analysis, Task Order: VA118A-11-0178, 2012.

The MITRE Corporation. "5.1.6 A&B: Classification Groups Research, Proposal, and Analysis." VA Disaster Recovery Analysis Project, Task Order:VA118A-11-0178, 2012.

The MITRE Corporation. "5.1.7 A: Public Cloud DR Viability Research Summary." VA Disaster Recovery Analysis, Task Order:VA118A-11-0178, 2012.

The MITRE Corporation. "5.1.7 B: Criteria for Public Cloud DR Viability ." VA Disaster Recovery Analysis, Task Order: VA118A-11-0178, 2012.

The MITRE Corporation. "5.1.9: Recommended DR Technology Solutions, Platform Definitions, and Solution Classifications." VA Disaster Recovery Analysis, Task Order: VA118A-11-0178, 2012.

The MITRE Corporation. "VA Enterprise Disaster Recovery Technology Solutions Standard, Version 1.0." VA Disaster Recovery Analysis, Task Order: VA118A-11-0178, 2012.

U.S. Chief Information Officer, Vivek Kundra. *25 Point Implementation Plan to Reform Federal Information Technology Management.* December 9, 2010.