admin

THE SOUTH CAROLINA
DEPARTMENT *of* ADMINISTRATION

South Carolina
Statewide Information Technology Shared Services

Enterprise Technology Architecture

RACI matrix for ETA

Suggested ETA Guidance Area	Executive Oversight Board	AWG	TWG	SARB	Agency Personnel (Architects)	DTO
Set Shared Services vision, strategy, and objectives	**R, A**	**C**	**C**	**I**	**I**	**I**
Define functional scope of new or changed shared services	**A**	**R**	**C**	**C**	**I**	**I**
Define technical requirements for new or changed shared services	**FA**	**C**	**A,R**	**C**	**C**	**C**
Establish Technology & Security Architecture - Patterns and Components	**FA**	**I**	**A**	**R**	**C**	**C**
Apply ETA to SSO projects			**I**	**I**	**C**	**A,R**
Apply ETA to Agency projects		**I**	**I**	**I**	**A,R**	**C**
Review Agency and DTO projects for ETA compliance	**I**	**I**	**A**	**R**	**C**	**C**
Grant Exceptions to ETA	**FA**	**C**	**A**	**R**	**I**	**I**
Adapt ETA to address Agency need		**I**	**FA**	**A**	**R**	**C**
Adapt ETA to address Shared Service need			**FA**	**A**	**C**	**R**
Define methodologies/tools/design patterns for SSO environment	**I**	**I**	**A**	**R**	**C**	**C**
Provide agency requirements for SSO capacity			**I**	**I**	**A,R**	**I**
Capacity Planning for SSO environment	**I**	**I**	**I**	**C**	**I**	**A,R**

FA = Final Approver R = Responsible A = Accountable/Approver
C = Consulted I = Informed

Enterprise Technology Architecture Principles – 1 of 2

Principle	Question to Answer	SC position
Time Horizon	What is our planning horizon for the architecture?	The state's Enterprise Technology Architecture will address a 3-year target state horizon
Architecture Scope	What environment(s) will be addressed by the ETA?	The Enterprise Technology Architecture will apply to all IT at all state agencies
Architecture Compliance	What degree of compliance is desired?	We seek 100% compliance with the ETA but will allow exceptions through a well-defined governance process
Architecture Diversity	Do we seek to establish one or many architectural models for each technology capability?	We strive to standardize on as few architectural approaches as possible, but expect that the number will vary by domain and subject area
Technology Maturity	What is our preference for the market maturity of our technologies?	The ETA will favor market-tested technologies, but will allow less mature technologies to be used where needed to respond to new risks, such as in Security Technology
Vendor Market Position	What is our preference for the market position of vendor/product standards, where needed?	Where vendors or products are identified in the ETA, we will favor vendors that are identified as Market Leaders or Challengers

Enterprise Technology Architecture Principles – 2 of 2

Principle	Question to Answer	SC position
Solution Integration	What is our preference for integrated product suites vs best of breed?	The state has a preference for solutions that minimize the need for custom integration, but expects high variability from one solution to another
Cost Sensitivity	What is the posture towards costs of infrastructure?	The Architecture should balance cost containment and investment in new technology
Proprietary & Open Standards	What is our position on open standards?	The ETA will favor vendor-agnostic standards but allow for vendor-specific architectural standards where required
Buy / Build	What is our preference for buy, build and reuse for infrastructure software	The ETA will specify services and packaged applications over custom developed software for any infrastructure applications
Technology Management	What degree of in-person IT support is preferred for infrastructure?	The ETA will strongly favor technologies that support high automation and limited management
Preservation of Current Architecture	What degree of change from the current state is appropriate when developing the ETA?	The ETA will leverage agency best practices wherever possible but will allow for new approaches and solutions

ETA Artifacts: Patterns

- Patterns are illustrations that show how technology components interact to deliver infrastructure functionality.
  - They provide specific guidance for use by solution designers and implementers on how technologies need to interoperate to provide capabilities to users and systems.
  - They do not represent detailed designs for solutions which must be determined based on specific business requirements.
  - They typically show technology categories but usually not products; site types but usually not specific locations; and levels of capacity/availability needed but usually not the number of units, speeds and feeds

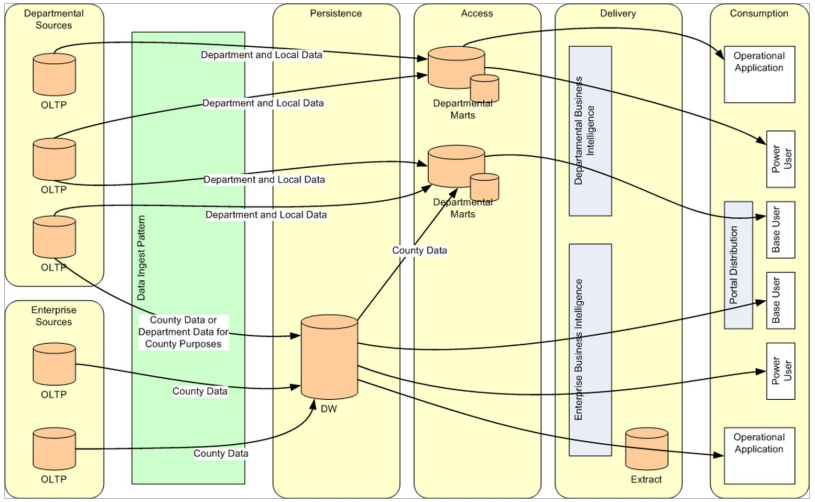

Data Management Pattern

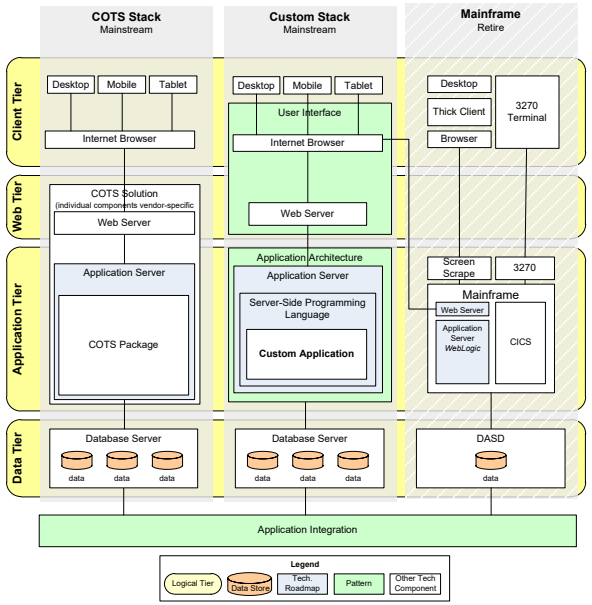

Application Platform Pattern

ETA Artifacts: Components

- Components specify the Enterprise standards for that technology, and organize them by lifecycle designation. Components provide a view of current technology standards and their evolution over time.
- Components also provide information to guide transition planning from current to target state

Portal Server	
Definition	**Implementation Guidance**
A portal is Web software that provides access to and interaction with relevant information assets by select targeted audiences, delivered in a highly personalized manner. Enterprise portals may face different audiences. Vertical portals focus on accessing specific applications or business functions. Horizontal portals (which is the subject of this brick) integrate and aggregate information from multiple cross-enterprise applications, as well as specific line-of-business tools and applications.	▪ Currently, Microsoft SharePoint is the standard, but is only used for internal/employee audiences ▪ Microsoft Dynamics and ADXStudio will be used, but these products do not excel as a Portal Server ▪ EMC Documentum Web Publishing currently used for external audiences, but is being retired. The replacement for external portals is still to be determined. There is an active project to define strategy for external Web presence ▪ WebParts are the only accepted portal-building technology. There are no Java portlets in the environment ▪ Adobe Live Cycle is used for forms management (with Kofax and Captiva for capture). It is only to be used as a portal for forms-related capabilities. ▪ DotNetNuke is used for an internal internet portal for CalWin, which must be migrated.
Retirement	**Mainstream**
DotNetNuke EMC Documentum Web Publishing Oracle Portal	Microsoft SharePoint 2010 Adobe Live Cycle
Containment	**Emerging**
ADX Studio Microsoft Dynamics CRM	Microsoft SharePoint 2013 Adobe CQ5 Oracle WebCenter

Typical Technology Lifecycle

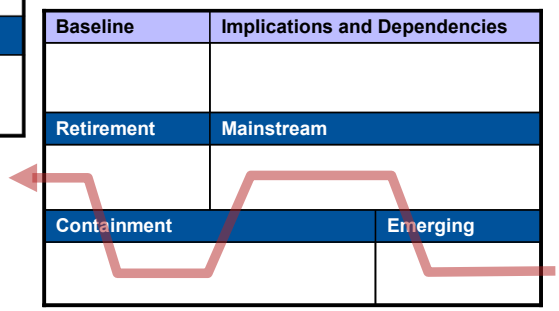

Table of Contents

Shared Services Data Center Pattern

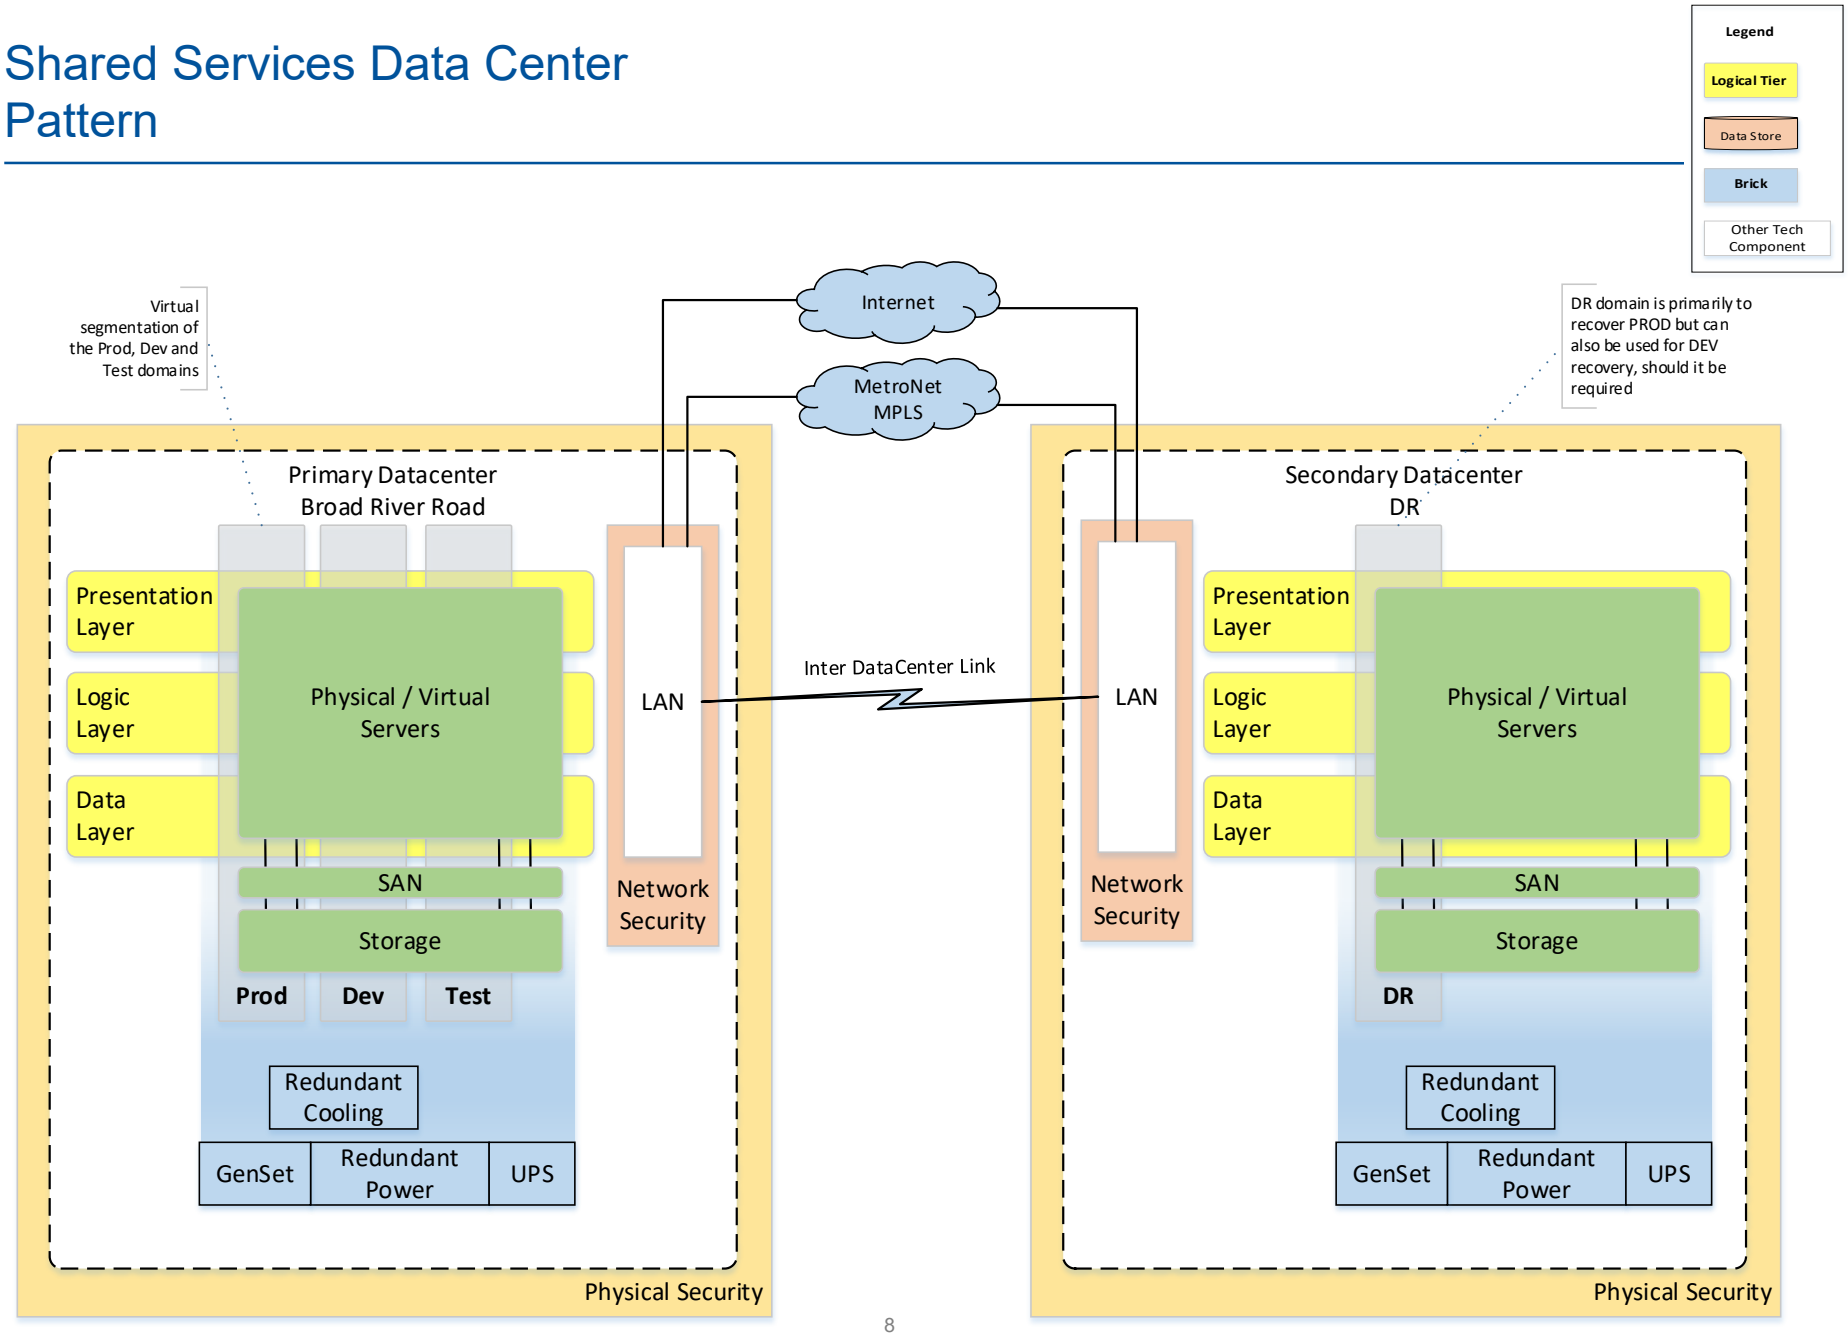

Shared Services Data Center Context

- **Location:**
  - Primary and Secondary datacenters approximately 100 – 115 miles apart
- **Power:**
  - Redundant Power (external feed from 2 distinct substations)
  - UPS (sized to maximum load for 5-10 minutes)
  - Gen Set (sized to maximum for 72 hours)
- **Cooling:**
  - Redundant Cooling (n+1 chilled water CRAC units)
- **Physical:**
  - 7x24 Operations (onsite)
  - 7x24 Protective Services [Primary Site only]
  - Dry pipes fire Suppression system [Primary Site] and Inert Gas fire Suppression system [Secondary]
  - 2-factor (badge + biometric) access control [Primary Site only]
- **Environments:**
  - Prod, Dev and Test at Primary Data Center
  - DR at Secondary Data Center
- **Networking:**
  - VLan-based segmentation by client and workload type
  - Access to external networks
    - Internet (at Primary and Secondary datacenters)
    - MetroNet MPLS (at Primary and Secondary datacenters)
    - Inter Data Center Link (Metro-E)
- **Data Center Management tool:**
  - None
  - Johnson Controls on cooling and power components

Shared Services Data Center
Technology Architecture Brick: Power

Shared Services Data Center - Power	
Baseline	**Guidance**
▪ Redundant Building Power Sources ▪ Uninterruptible Power Supply (UPS) / Batteries ▪ Diesel Generator ▪ Redundant Intra-Data-Center Power Distribution and cabling to each equipment rack	▪ Ensure redundant Power feeds into the Data Center, preferably from two substations, using different ingress paths (north and south routes) ▪ Leverage double-conversion UPS technology to condition power before it reaches equipment within the datacenter. UPS batteries should be able to maintain power for 5 to 10 minutes to allow for generator start-up ▪ Test generator twice yearly and ensure regular yearly maintenance of generator parts and fuel supply. Keep at least 48 hours of generator fuel on hand. ▪ Employ rigorous capacity management to stay ahead of equipment demands. Include this exercise as part of change management for infrastructure.
Retirement	**Mainstream**
▪ N/A	▪ Redundant Building Power Sources ▪ Uninterruptible Power Supply (UPS) / Batteries ▪ Diesel Generator ▪ Redundant Intra-Data-Center Power Distribution and cabling to each equipment rack
Containment	**Emerging**
▪ N/A	▪ Modular self-contained Aisle and Racking System

Shared Services Data Center
Technology Architecture Brick: Cooling

Shared Services Data Center - Cooling	
Baseline	**Guidance**
▪ Redundant Chilled Water CRAC Units	▪ Use a n+1 approach to overprovisioning cooling capacity to allow for CRAC unit maintenance or failure. ▪ Employ rigorous capacity management to stay ahead of equipment demands. Include this exercise as part of change management for infrastructure. ▪ Investigate approaches to optimize cooling and airflow within the datacenter in order to maximize energy efficiency.
Retirement	**Mainstream**
▪ N/A	▪ Redundant Chilled Water CRAC Units
Containment	**Emerging**
▪ N/A	▪ Hot/Cold Aisles containment ▪ Modular self-contained Aisle and Racking System

Shared Services Data Center
Technology Architecture Brick: Data Center Network Connectivity

Shared Services Data Center – Data Center Network Connectivity	
Baseline	**Guidance**
▪ Redundant Internet connectivity [Primary Site only] ▪ Redundant Metro-Net (MPLS) connectivity [Primary Site only] ▪ Redundant Inter Data Center (Primary-Secondary) connectivity over Private links	▪ Where possible, source redundant connectivity through different providers, using distinct ingress points and routes. ▪ Leverage Metro-Net for client access where technically and economically feasible, otherwise use VPN over public Internet. Consider wireless (4G or Satellite) backup to minimize the impact of WAN outages. ▪ Select a link technology that will minimize latency (e.g. dark fibre, private links) for Intra Data Center communication to allow for extended clustering technology as well as low Recovery Point Objectives (RPO).
Retirement	**Mainstream**
▪ N/A	▪ Redundant Internet connectivity [Primary and Secondary sites] ▪ Redundant Metro-Net (MPLS) connectivity [Primary and Secondary sites] ▪ Redundant Inter Data Center (Primary-Secondary) connectivity over Private links
Containment	**Emerging**
▪ N/A	

Enterprise Storage Pattern

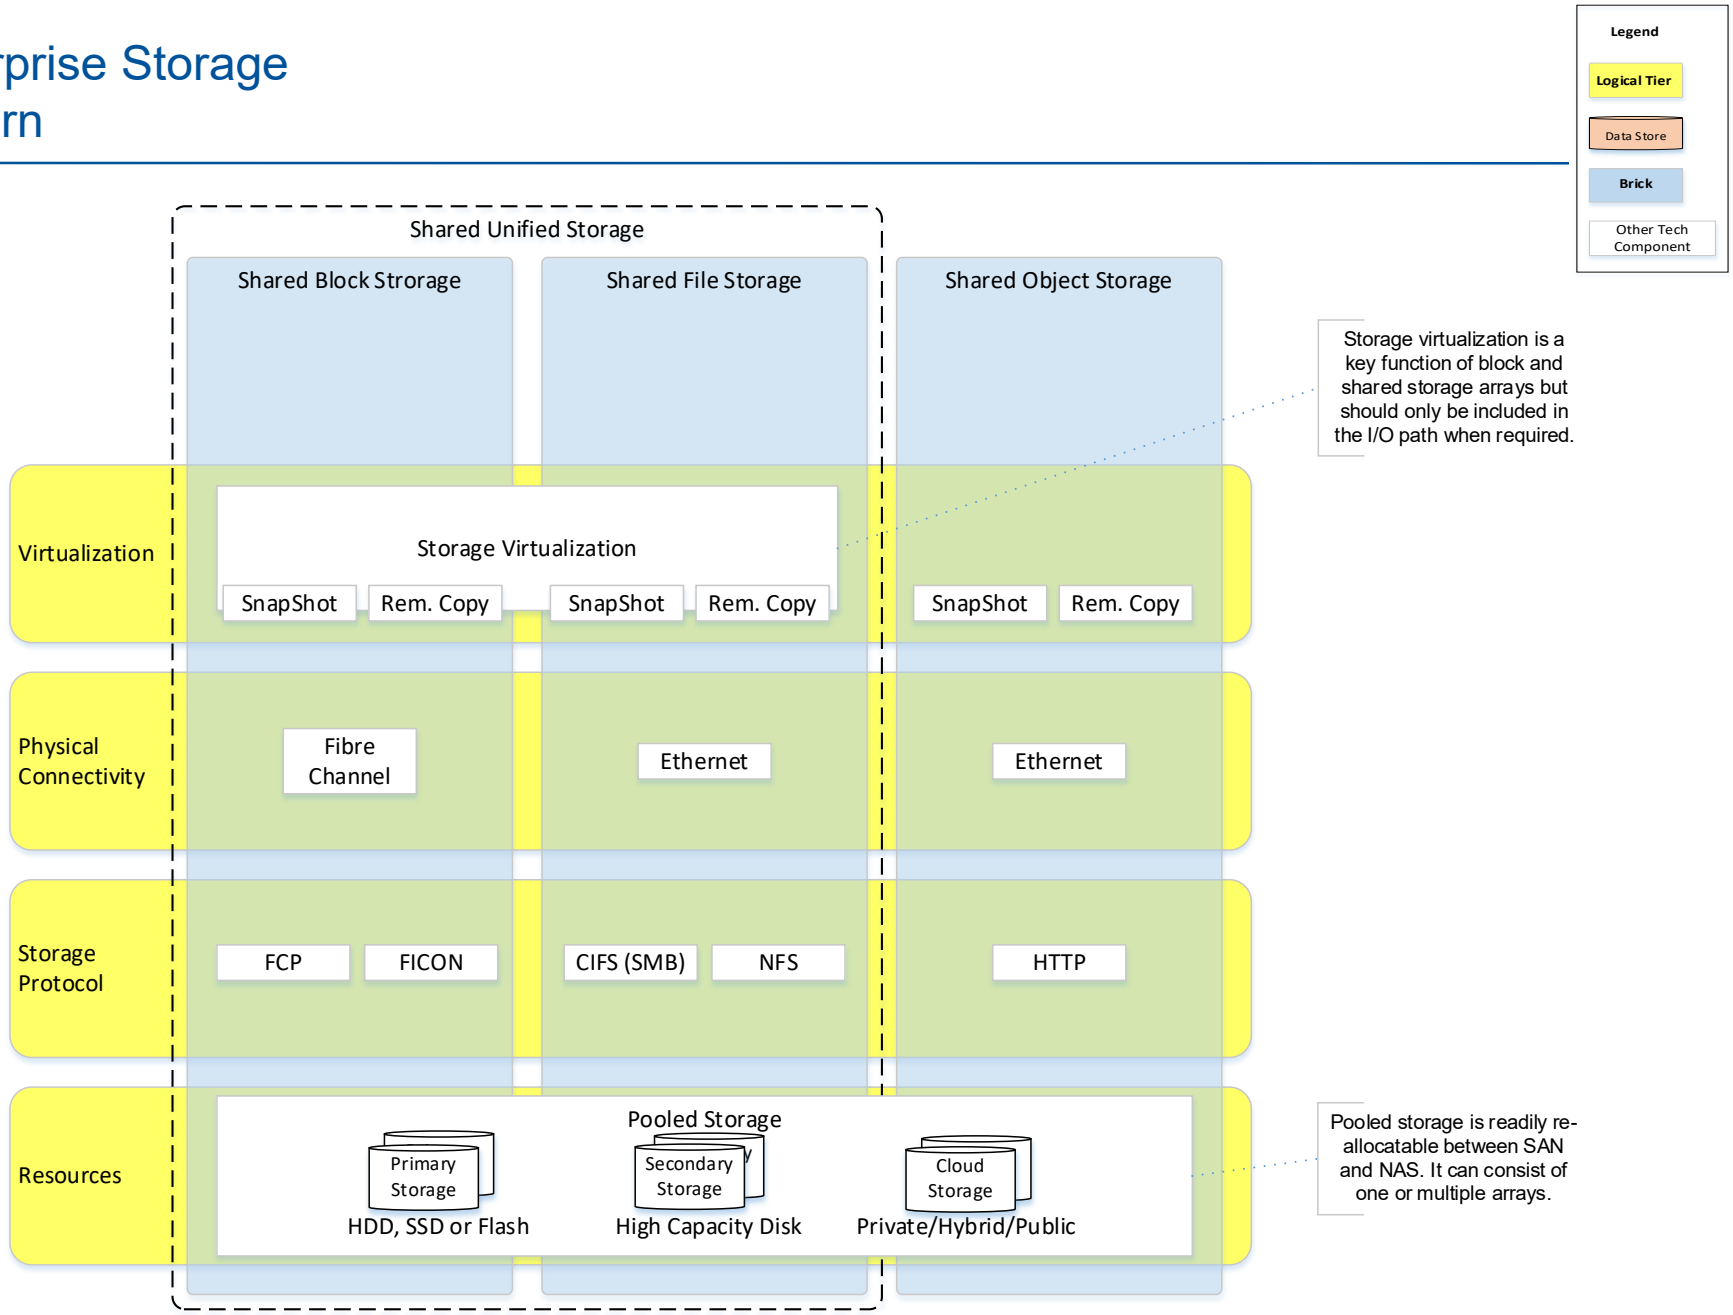

Enterprise Storage Context

- **Storage Architecture Type:**
  - Shared File Storage
    - Hitachi NAS Platform
    - Multiple department and agencies are using NetApp and rely on advanced features (e.g. SnapVault, SnapManager for SQL) for data protection
  - Shared Block Storage
    - Hitachi G1500 for Open Systems
    - IBM DS8870 for Mainframe
    - A variety of storage platforms are used by department and agencies
  - Shared Object Storage
    - Hitachi HCP
- **Storage Virtualization**
  - HDS G1500 can provide storage virtualization
- **Storage Management**
  - DTO, as well as the departments and agencies surveyed use native element managers to manage storage

Enterprise Storage
Technology Architecture Brick: Shared Block Storage

Enterprise Storage - Shared Block Storage	
Baseline	**Guidance**
▪ HDS G1500 for Open Systems Storage ▪ IBM DS8870 for Mainframe storage	▪ While there is a logical division between open systems and mainframe storage, there might be efficiency gained from consolidating both workloads in a single footprint. ▪ Storage virtualization (either built into the storage system or as an appliance that is external to the storage system) will simplify storage management and data migration. HDS' G1500 platform provides this functionality; care should be taken to maintain it if the G1500 is ever replaced. ▪ NetApp technology is used as Shared Unified Storage by multiple departments and agencies. In some cases, backup and DR processes leverage NetApp specific features and APIs such as SnapVault or SnapManager for SQL, which can often be 'ported' to a different storage array but may require a change in scripts or even applications. ▪ At a minimum, the Shared Block Storage platform should offer Snapshot and Remote Copy capability. VM has this not needed. ▪ HDS, IBM and NetApp storage can be found in the leader's quadrant of Gartner's Magic Quadrant for General-Purpose Disk Arrays
Retirement	**Mainstream**
▪ Non-enterprise storage arrays such as HP EVA and Dell Compellent ▪ Direct attached storage ▪ All other non mainstream or containment arrays	▪ Hitachi (for Open Systems), including storage virtualization ▪ IBM (for Mainframe)
Containment	**Emerging**
	▪ Cloud based storage ▪ All-flash storage arrays ▪ Software-Defined storage

Enterprise Storage
Technology Architecture Brick: Shared File Storage

Enterprise Storage – Shared File Storage	
Baseline	**Guidance**
▪ Hitachi NAS	▪ At Scale, NAS workload should reside on purpose-built arrays rather than being served from a Windows or Linux host front-ending block storage ▪ Look for NAS arrays that can be managed centrally and allow for workload migration between arrays for load balancing. ▪ Favor Unified Storage arrays that provide both block and file storage within the same array, in order to minimize management overhead and maximize storage efficiency. ▪ At a minimum, the Shared File Storage platform should offer Snapshot and Remote Copy capability. ▪ NAS arrays should include replication
Retirement	**Mainstream**
▪ Windows or Linux file servers ▪ End-of-life NAS arrays, such as EMC Celerra ▪ All other non-mainstream arrays	▪ Hitachi NAS
Containment	**Emerging**
	▪ Scale-out distributed file system, with global namespace

Enterprise Storage
Technology Architecture Brick: Shared Object Storage

<table>
<tr><th colspan="3">Enterprise Storage – Shared Object Storage</th></tr>
<tr><th>Baseline</th><th colspan="2">Guidance</th></tr>
<tr><td>▪ Hitachi HCP</td><td colspan="2">▪ Infrastructure and operations groups are attracted by the lower total cost of ownership and the scalability of object storage, whereas enterprise developers are attracted to its programmability, cloud portability and productivity improvements, both contributing to strong growth of this segment
▪ Choose object storage products as alternatives to block and file storage when you need huge scalable capacity, reduced management overhead and lower cost of ownership
▪ Build on-premises object storage repositories with the hybrid cloud in mind, and evaluate their API support and level of compatibility with dominant public cloud providers for data portability
▪ At a minimum, the Shared Object Storage platform should offer Snapshot and Remote Copy capability.
▪ Hitachi HCP is found in the challenger quadrant of the Gartner Magic Quadrant for Distributed File Systems and Object Storage</td></tr>
<tr><th>Retirement</th><th colspan="2">Mainstream</th></tr>
<tr><td>▪ End-of-life Object Storage, such as EMC Centera
▪ All other non mainstream or containment arrays</td><td colspan="2">▪ Hitachi HCP</td></tr>
<tr><th colspan="2">Containment</th><th>Emerging</th></tr>
<tr><td colspan="2">▪ N/A</td><td>▪ Cloud Storage Gateway</td></tr>
</table>

Storage Area Network Pattern

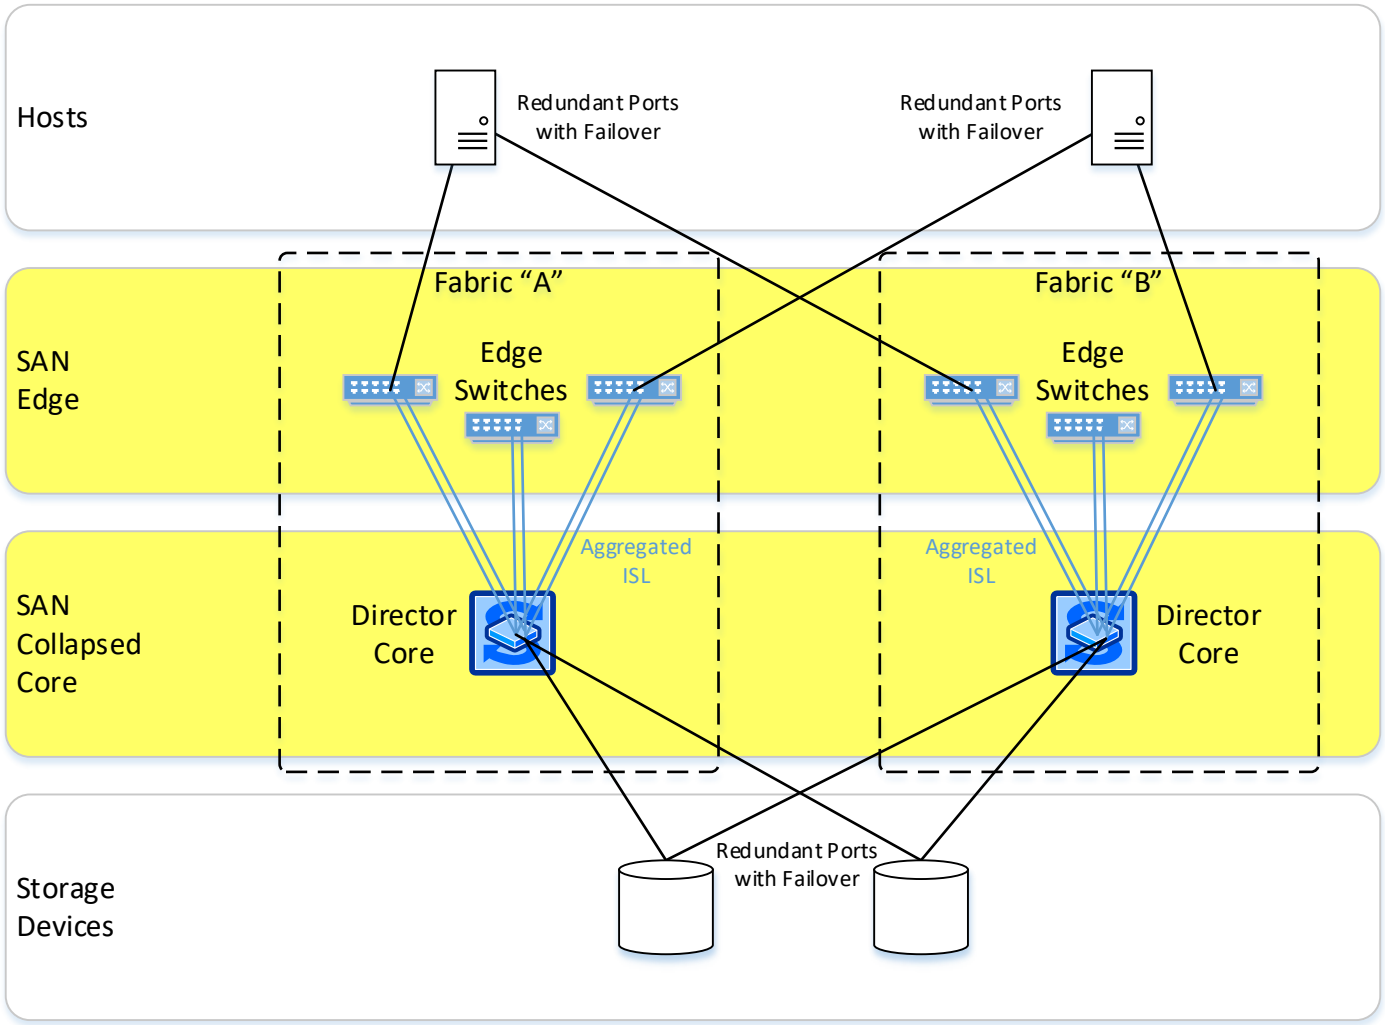

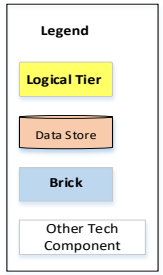

Storage Area Network Context

- **Storage Area Network Type:**
  - Based on Cisco MDS technology
    - Cisco MDS 9710 Directors currently in use at the Primary site
    - MDS 9250i switches currently in use at the secondary site
  - Mainframe SAN based on IBM (Brocade) technology
- **Storage Area Network Architecture:**
  - Redundant fabrics
  - Collapsed Core architecture:
    - UCS chassis with integrated edge switches
    - Storage connected directly to the Core
- **Open Systems**
  - Fibre Channel Protocol (FCP)
- **Mainframe**
  - FICON (over FC)

Storage Area Network
Technology Architecture Brick: Switches and Directors

<table>
<tr><th colspan="3">Storage Area Network (SAN) – Switches and Directors</th></tr>
<tr><th>Baseline</th><th colspan="2">Guidance</th></tr>
<tr><td>▪ Cisco 9710 directors (core) and 6248 switches (edge – part of UCS)
▪ Cisco 9250i switches
▪ IBM (Brocade) FICON directors</td><td colspan="2">▪ Single-vendor SAN are recommended as mixing technologies from multiple vendors will lead to reduced feature sets (compatibility mode) and may be difficult to support.
▪ This includes extended links (e.g. to a remote datacenter) as well; DTO should stick to one SAN vendor/technology through its environment.</td></tr>
<tr><th>Retirement</th><th colspan="2">Mainstream</th></tr>
<tr><td>▪ Any non mainstream or containment SAN technology</td><td colspan="2">▪ Cisco 9710 directors (core) with 16Gb/s Fibre Channel SFPs
▪ Cisco 6248 edge switches with 16Gb/s Fibre Channel SFPs
▪ Cisco 9250i switches
▪ IBM (Brocade) FICON directors for the mainframe environment only</td></tr>
<tr><th colspan="2">Containment</th><th>Emerging</th></tr>
<tr><td colspan="2">▪ Switches and Directors supporting less than 10Gb/s link speed</td><td>▪ 32 Gb/s FC</td></tr>
</table>

Backup, Restore and Recovery Pattern

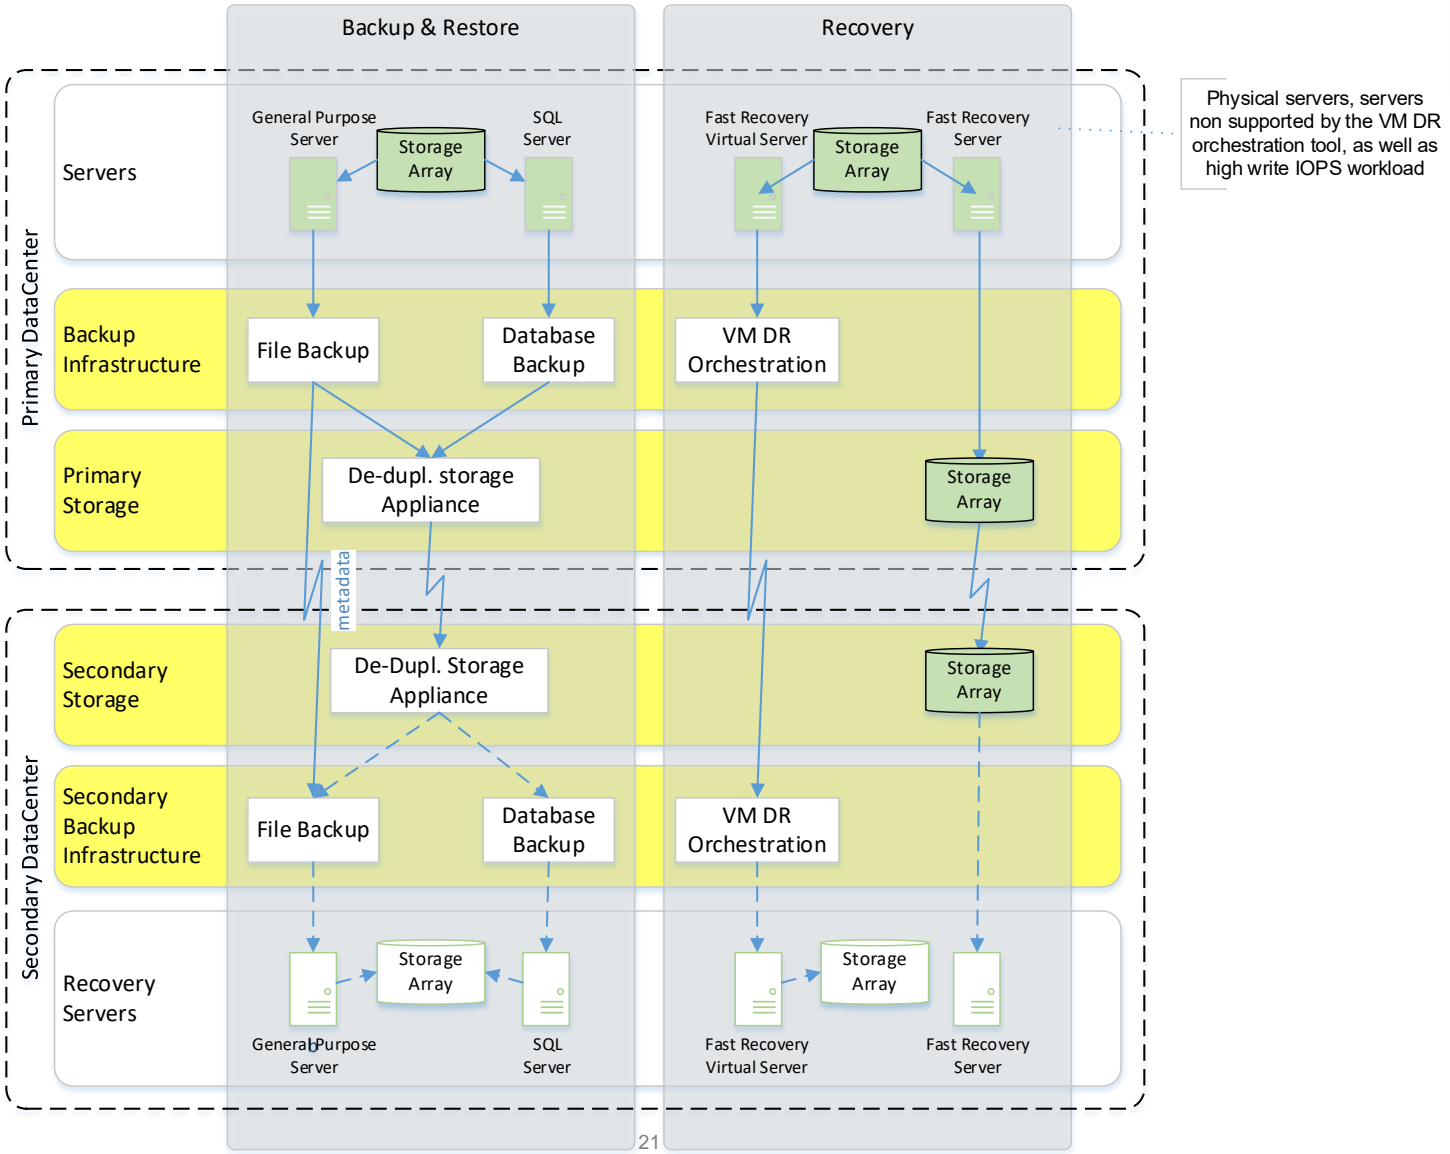

Backup, Restore and Recovery Context

- **Data Protection:**
  - Tools for backup (data protection):
    - Avamar + DataDomain for files
    - Idera + DataDomain for SQL Databases

- **Data Replication**
  - two approaches for DR (Replication and Recovery):
    - Snapshot + Remote Replication (on HDS)
    - VMWare SRM / Zerto

Backup, Restore and Recovery
Technology Architecture Brick: Backup and Restore

<table>
<tr><th colspan="3">Backup, Restore and Recovery – Backup and Restore</th></tr>
<tr><th>Baseline</th><th colspan="2">Guidance</th></tr>
<tr><td>▪ Avamar for backing up files
▪ Idera for backing up SQL databases
▪ DataDomain as back-end storage for backup services</td><td colspan="2">▪ Both Avamar and Idera integrate well with DataDomain, enabling appliance-based de-duplication to take place and minimize the size of backups
▪ One of Avamar’s strength is backing up remote servers over WAN links as it will employ link optimization by de-duplicating information sent over the WAN. DTO could potentially offer this service to agencies and departments that have servers outside of DTO’s datacenters.
▪ Idera SQL Safe backup optimizes the online backup of SQL servers by employing dynamic compression. It also provides “instant restore”, which brings databases back online immediately, without requiring a long restore process.</td></tr>
<tr><th>Retirement</th><th colspan="2">Mainstream</th></tr>
<tr><td>▪ Tivoli Storage Manager
▪ EMC Networker</td><td colspan="2">▪ EMC Avamar
▪ Idera SQL Safe backup
▪ DataDomain</td></tr>
<tr><th colspan="2">Containment</th><th>Emerging</th></tr>
<tr><td colspan="2">▪ N/A</td><td>▪ Archive Platform: (Data Archiving Solution)</td></tr>
</table>

Backup, Restore and Recovery
Technology Architecture Brick: Replication and Recovery

<table>
<tr><th colspan="3">Backup, Restore and Recovery – Replication and Recovery</th></tr>
<tr><th>Baseline</th><th colspan="2">Guidance</th></tr>
<tr><td>▪ VMWare Site Recovery Manager (SRM)
▪ Zerto
▪ Hitachi Universal Replicator</td><td colspan="2">▪ Virtual Machines: Both VMWare SRM and Zerto offer virtual machine replication across datacenters. These solutions ensure that virtual machines can be quickly ‘restarted’ at the secondary site, in the event that the protected virtual machine at the primary site becomes unavailable.
▪ Data Replication: Data replication allows for any data to be replicated across datacenters. This replication can be done synchronously (zero lag) using Hitachi TrueCopy using Hitachi Universal Replicator at any distance. Consistency points, usually created by using array-based snapshotting technology, are recommended, in order to ensure a recoverable copy is present at the secondary site.
▪ Data Replication: File and content systems should be capable of performing snapshots and replicating data to secondary systems.</td></tr>
<tr><th>Retirement</th><th colspan="2">Mainstream</th></tr>
<tr><td>▪ VMWare Site Recovery Manager (SRM)
▪ Hitachi Universal Replicator</td><td colspan="2">▪ Zerto for virtual servers that require quick recovery</td></tr>
<tr><th colspan="2">Containment</th><th>Emerging</th></tr>
<tr><td colspan="2">▪ N/A</td><td>▪ Cloud-based recovery service provider
▪ Container Management
▪ Long-distance Live VM migration</td></tr>
</table>

Enterprise Server Component

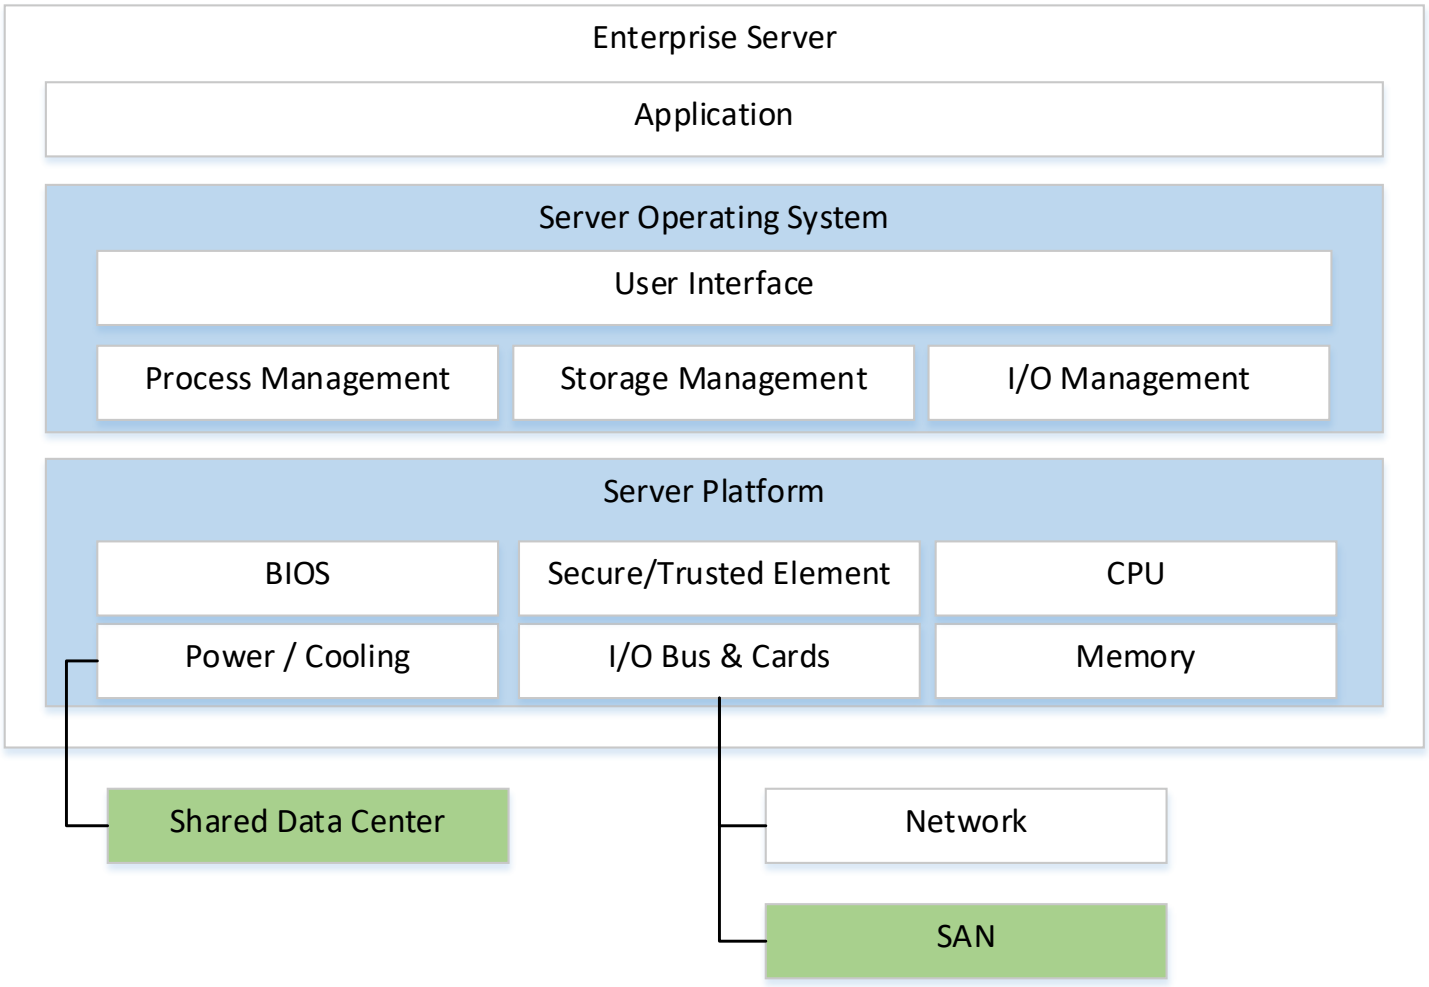

Enterprise Server Context

- **Compute platforms:**
  - Main compute platform is x86
    - Leveraging Cisco's UCS Integrated Systems platform to host VMWare, providing horizontal scalability in a blade form factor
    - Some legacy HPE ProLiant servers are also present but this form factor is no longer favored
  - There is a small number of RISC/Unix Servers
    - IBM POWER
  - DTO hosts a mainframe environment for several department and agencies, under a multi-tenancy model
    - Based on IBM z technology, using LPARs to segregate clients workloads

- **Operating Systems:**
  - On x86
    - Windows
    - Linux
  - On IBM POWER
    - AIX
  - On IBM z
    - zOS

Enterprise Server
Technology Architecture Brick: Server Hardware

<table>
<tr><th colspan="3">Enterprise Server – Server Hardware</th></tr>
<tr><th>Baseline</th><th colspan="2">Implications and Dependencies</th></tr>
<tr><td>▪ Cisco UCS (Unified Computing Systems)
▪ HPE ProLiant, Dell, IBM/Lenovo x86 servers
▪ IBM z (Mainframe)
▪ IBM POWER (RISC)
▪ Sun/Oracle
▪ HPE
▪ IBM i (AS/400)
▪ Unisys Mainframe</td><td colspan="2">▪ Through its integrated architecture and validated designs, UCS simplifies acquisition, scalability and support of DTO’s x86 server environment.
▪ Where possible (if the software and non-functional requirements allow it), the industry is shifting away from proprietary server stacks such as IBM i, Unisys and IBM z Mainframes towards commodity x86 servers. This usually leads to lower hardware and maintenance costs but may require increased focus on availability, manageability and scalability mechanisms to be able to match what has been offered by proprietary platforms.</td></tr>
<tr><th>Retirement</th><th colspan="2">Mainstream</th></tr>
<tr><td>▪ HPE ProLiant x86 rack mounted Servers
▪ IBM POWER (RISC)
▪ Sun/Oracle
▪ HPE
▪ IBM i (AS/400)
▪ Server appliances (physical)</td><td colspan="2">▪ Cisco UCS
▪ Virtual server appliances</td></tr>
<tr><th colspan="2">Containment</th><th>Emerging</th></tr>
<tr><td colspan="2">▪ IBM z (Mainframe)
▪ Unisys Mainframe</td><td>▪ Hyperconverged Integrated Systems</td></tr>
</table>

Enterprise Server
Technology Architecture Brick: Server Operating Systems

<table>
<tr><th colspan="3">Enterprise Server – Server Operating Systems</th></tr>
<tr><th>Baseline</th><th colspan="2">Implications and Dependencies</th></tr>
<tr><td>▪ x86 Servers
▪ Microsoft Windows Server
▪ RedHat Enterprise Linux, SUSE, CentOS
▪ Mainframe
▪ IBM zOS
▪ Unix:
▪ IBM AIX
▪ Solaris
▪ HP-UX
▪ Other:
▪ IBM i (OS/400)</td><td colspan="2">▪ SSO current strategic direction aligns with Gartner best practices. SSO should continue to proactively contain and retire RISC/Unix systems and actively support businesses in the migration from RISC/Unix and other proprietary server operating systems to x86 architecture where and when possible.
▪ Gartner is forecasting a more rapid decline in the RISC/Unix marketplace over the next 5 years resulting in materially reduced ISV product support and availability.
▪ The x86 platform continues to grow while proprietary Server Operating Systems continue to decline. For example, in the 2015 server OS market, Windows (server) and Linux (server) are the only two subsegments that grew positively at 7.7% and 12.4%, respectively. The combination of these two OSs has constituted a market share of 67.4% (although Linux being smaller) that demonstrates the prevalence of x86 platform in the world's server computing environment</td></tr>
<tr><th>Retirement</th><th colspan="2">Mainstream</th></tr>
<tr><td>▪ Microsoft Windows Server, unsupported versions of OS
▪ All Linux distributions other than RedHat Enterprise Linux, unsupported versions of RedHat Enterprise Linux
▪ All Unix OS, including IBM AIX, Solaris and HP-UX
▪ All other OS, including OS/400 and Unisys</td><td colspan="2">▪ Microsoft Windows Server, Latest release
▪ Red Hat Enterprise Linux, Latest release</td></tr>
<tr><th colspan="2">Containment</th><th>Emerging</th></tr>
<tr><td colspan="2">▪ IBM zOS
▪ Microsoft Windows Server, Still Supported OS version, not latest.
▪ RedHat Enterprise Linux, , Still Supported OS version, not latest.
▪ Unisys OS</td><td></td></tr>
</table>

Server Virtualization Pattern

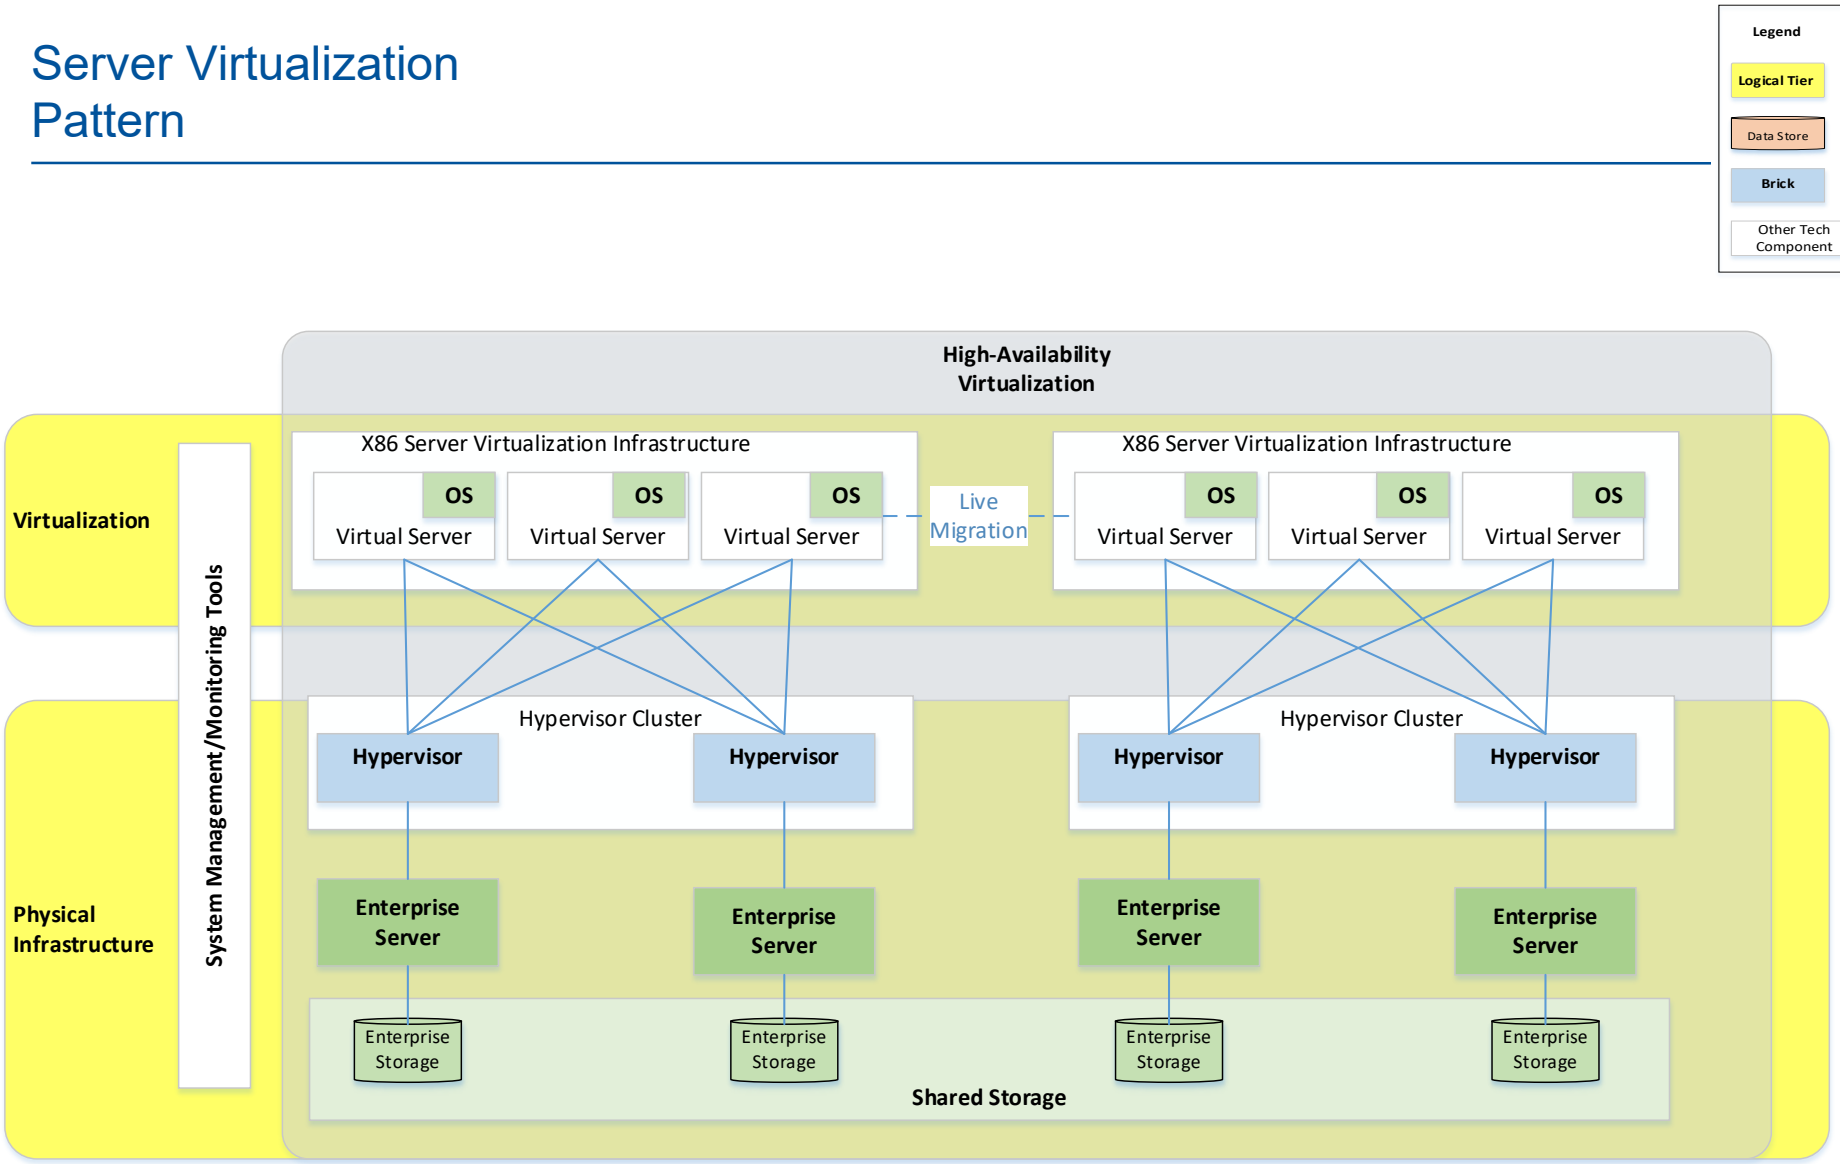

Server Virtualization
Description

x86 server virtualization enables multiple operating system (OS) instances to be deployed, operated and managed concurrently on a single physical x86 server. x86 server virtualization is enabled by a virtual machine monitor (VMM), which runs on top of a hypervisor (most common) layer.

Virtual machines (VMs) should be hosted on standardized configuration of blade servers or rack mounted x86 servers. Currently, industry norm ranges between 4-8 VMs per core. Standard file based and host based backup solutions should be in place to manage backup of file systems and data. Images and data should be hosted in a SAN environment for ease of management and migration (if needed). Separate test/development and production environments where possible.

Virtual stack management in terms of hardware management, VM management, capacity planning, workload planning and measurement, and various utility tools that are coming up in the market should be provisioned to provide management of the environments.

Although virtualization is applicable to the vast majority of workloads, not all are easy to virtualize as VMs. Workloads that require high CPU utilization, high memory and higher amounts of Input/Output may not be ideal candidates for virtualization. VM workloads should be selected based on a virtual strategy and workload analysis prior to hosting on virtual environments.

Server Virtualization Context

Hypervisor:

- DTO is currently using VMWare for x86 Server virtualization
- Most department/agencies also use VMWare while some use Hyper-V (mostly in smaller deployment)

VM Workload Management:

- VCenter/VMotion is used to non-disruptively migrate workload from one server or UCS frame to another for balancing workloads and to minimize service disruption due to planned updates, upgrades or repairs
- Dynamic Resource Scheduler (DRS) is used to automatically migrate workload to balance resource utilization

Server Virtualization
Technology Architecture Brick: x86 Hypervisor

<table>
<tr><th colspan="3">Server Virtualization – x86 Hypervisor</th></tr>
<tr><th>Baseline</th><th colspan="2">Implications and Dependencies</th></tr>
<tr><td>▪ VMware vSphere
▪ Microsoft Hyper-V</td><td colspan="2">▪ Based on the understanding of the environment and the input reflected by stakeholders interviews, it is recommended that SSO maintain a single hypervisor for production workloads. However, it is understood that specialized heterogeneous virtualization solutions may be required in some instances due to technical, financial or political considerations (and these exceptions should be reviewed on a case by case basis)
▪ Microsoft Hyper-V has been used by a few agencies and has been deployed on select AD Domain Controller (Tier 0) which Microsoft is supporting. In the near term, Microsoft Hyper-V should be planned to be migrated to VMware in the future</td></tr>
<tr><th>Retirement</th><th colspan="2">Mainstream</th></tr>
<tr><td>▪ N/A</td><td colspan="2">▪ VMware vSphere</td></tr>
<tr><th colspan="2">Containment</th><th>Emerging</th></tr>
<tr><td colspan="2">▪ Microsoft Hyper-V</td><td></td></tr>
</table>

Client Virtualization Pattern

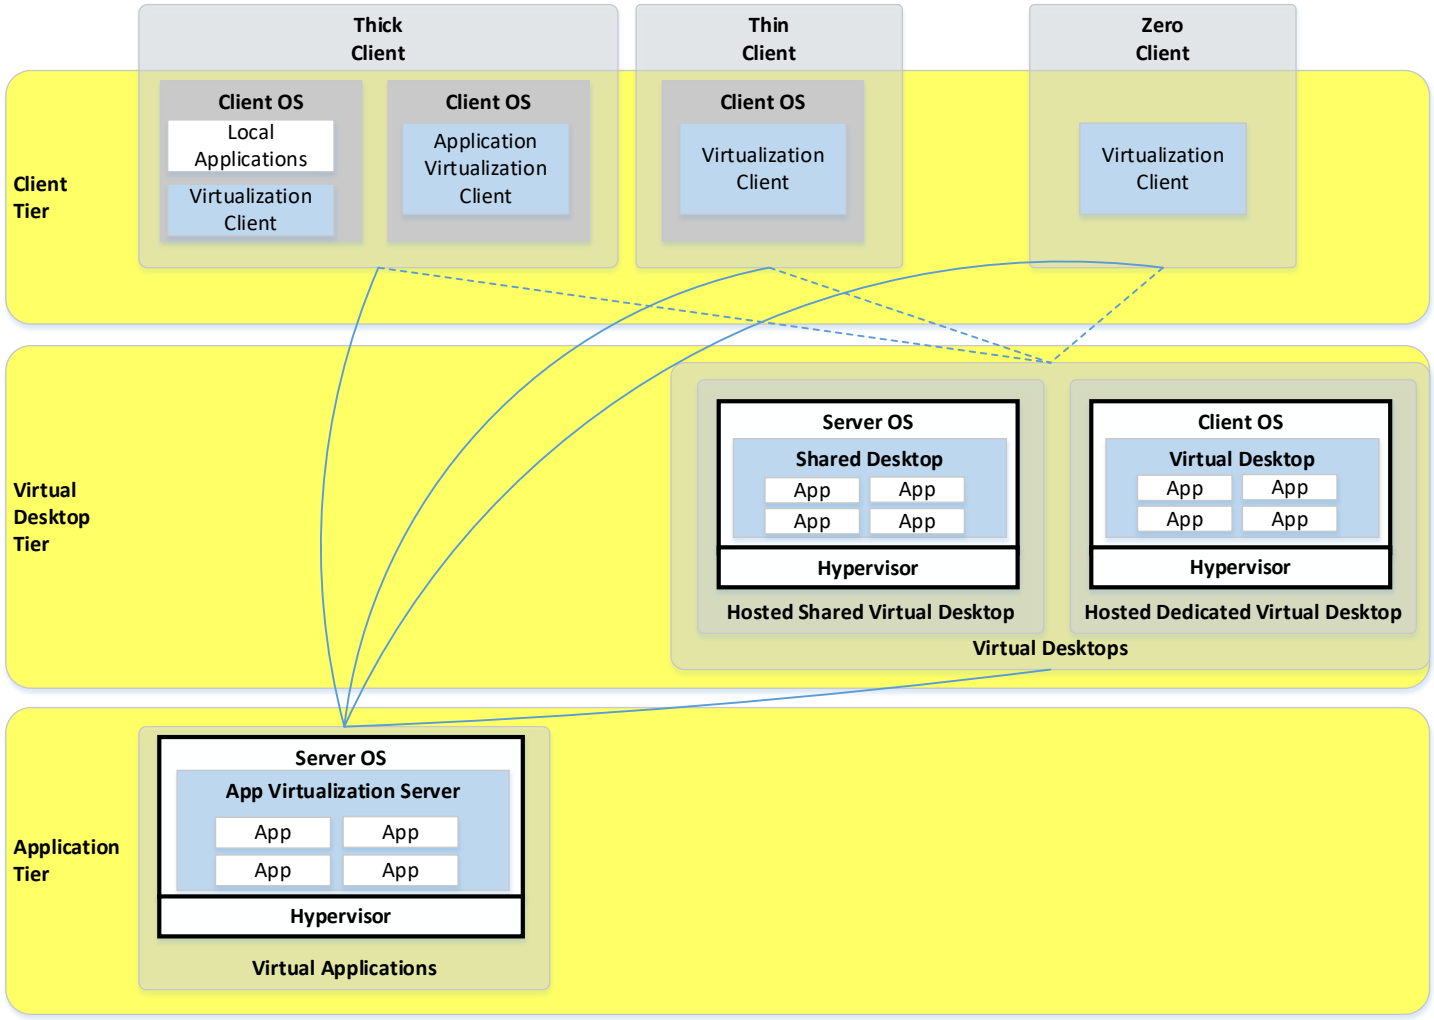

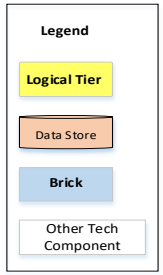

Client Virtualization
Description

For state of South Carolina- Shared Services Organization, client-side virtualization is defined by two capabilities—application and desktop—deployed across three end-point client types—thick, thin, and zero:

- **Application Virtualization** : Streamline PC application deployment, and address packaging and application coexistence problems.
- **Hosted Virtual Desktop**: Provide end-user flexibility, efficiency, possible energy savings and other benefits, enabling administrators to manage desktops from a centralized location and end users to access their desktops from machines in any location.
- **Thick Client**: A thick-client is generally considered to be a standard PC (in any of the various form factors such as desktop, laptop, and tablet) designed to run an off-the-shelf operating system such as Microsoft Windows.
- **Thin Client**: A thin-client is a purpose-built device with a hardware configuration and operating system that have been created specifically to fill the virtualized client role. The goal is to create a device that is simpler to manage, more secure, and more power efficient than a thick-client.
- **Zero Client**: A zero-client is built around a proprietary hardware platform, usually in the form of application-specific integrated circuit (ASIC), which includes processing power, memory, network, and graphics circuitry in a single chip. Zero-clients also eschew a commercial operating system platform such as Windows in favor of an embedded operating system.

Application virtualization with thick clients is mainstream at most agencies within the state, notably with the virtualization of the key GIS applications. This is delivered by an application virtualization server that delivers applications remotely to a thick client.

- Application virtualization should be used to virtualize non-standard desktop applications. Unique applications such as GIS or imaging systems that require heavy processing, or those that cause conflicts with other installed applications are candidates for virtualization. Microsoft Office and other standard desktop builds should still be installed on the client until sufficient SSO staff and budget are allocated to provide a comprehensive application virtualization strategy.
- However, application virtualization can also be used to virtualize all desktop applications for small environments with a single desktop build, and in which business users perform the same tasks and require little, if any, customization (e.g. call centers, administrative).
- Another common use case for desktop and application virtualization is environments that require additional security, particularly in scenarios where information exfiltration would have devastating consequences.
- Connectivity is important for application virtualization to be viable.

Client Virtualization Context

Desktop Virtualization:

- DTO is currently using VMWare Horizon
- Some agencies/departments use Citrix XenDesktop for Desktop virtualization, in addition to VMWare Horizon

Application Virtualization:

- DTO uses Citrix XenApps for Application virtualization

Major uses:

- Implementation in place for a few hundred endpoints
- Primary use today is for remote access, to increase security (remote endpoint may not be fully managed by agency/department or DTO)

Client Virtualization
Technology Architecture Brick: Application and Desktop Virtualization

<table>
<tr><th colspan="3">Client Virtualization – Application and Desktop Virtualization</th></tr>
<tr><th>Baseline</th><th colspan="2">Implications and Dependencies</th></tr>
<tr><td>Desktop Virtualization:
▪ VMware Horizon
Application Virtualization:
▪ XenApp</td><td colspan="2">▪ VMware Horizon is the mainstream product for desktop virtualization. VMware Thin App should be evaluated for deployments with both virtual applications and VMware desktops in the future
▪ There are a few instances where Citrix is used which should be contained to specific instances and should be migrated to VMware in the future.</td></tr>
<tr><th>Retirement</th><th colspan="2">Mainstream</th></tr>
<tr><td>▪ None identified</td><td colspan="2">Desktop Virtualization:
▪ VMWare Horizon
Application Virtualization:
▪ XenApp</td></tr>
<tr><th colspan="2">Containment</th><th>Emerging</th></tr>
<tr><td colspan="2">▪ Citrix XenDesktop
▪ ThinApp</td><td>▪ Application Virtualization: Microsoft Software Grid (Softgrid)</td></tr>
</table>