Coast Colleges

COASTLINE COMMUNITY COLLEGE

GOLDEN WEST COLLEGE GWC HUNTINGTON BEACH, CA

ORANGE COAST COLLEGE

Golden West College Server and Storage Consolidation Project

12/11/2014

Matt Marino, Kevin Masui, Danny Gorman, Eng Soedjono, Rich Filakousky

Executive Summary

- **Current MDF at GWC which houses server and storage is scheduled for demolition in December 2015**
- **Consolidate ageing End of Life (EOL) hardware to central management location at District Data Center (DC)**
- **Reduce Operating Expenditures and Capital Expenditures**
  - No need for 3rd party expensive warranty (**60k plus per year**)
  - Less power consumed and less cooling needed (ref slide 25)
- **Utilize latest hardware technology to consolidate and standardize GWC servers and storage**
  - 68 total servers consolidated to 4 physical servers (**94% reduction**)
- **Ability to provide security, redundancy, performance, scalability, and management**
  - Ability to provide best in class technology to students, faculty, and staff
  - N+1 redundancy within blades, chassis, network, and storage
    - Provide High Availability (**HA**) and Fault Tolerance (**FT**) to GWC infrastructure
    - Located in District DC for advanced power and cooling requirements
    - Infrastructure tuned to work towards 99.99% (52 mins of downtime per year) uptime compared to sub 99% (4 days per year) currently
  - Central virtual and hardware management via one central management console, Single Pane of Glass (**SPoG**)

Requirements

- Consolidate all server and storage hardware at District DC (Data Center) including the Virtual Desktop Infrastructure (VDI)
- Physical servers converted to virtualized servers
- Sufficient bandwidth and transfer windows to migrate from GWC to District DC over the WAN
- WSUS Patching to desktops over WAN - System
- Redundant WAN connectivity
- Increased storage space and performance
- Applications and owners of servers identified and communicated with
- Network at district is configured for migration
- Scheduled server downtime for migration of some resources

Risks

- Insufficient bandwidth between GWC and District can lead users to experience application and/or VDI latency
  - Coastline College currently functions over same WAN connection as proposed for GWC
- WSUS patching over the WAN may not be completed during the maintenance period carrying over to work hours. This will flood the network and create latency for end users.
  - Currently researching possible solutions, one being physical server onsite at GWC in new MDF for local security patching
- Loss of WAN connectivity will disable all services at GWC
  - Currently GWC relies on WAN for email, Banner, and other district hosted resources. Coastline operates with this model currently and experiences less than .067% of packet loss and averages 4% utilization of bandwidth (all server traffic and all VDI traffic for classes)
- Converting servers to virtual (**P2V – physical to virtual**) has the potential to break the services offered (i.e. SQL dependent applications)
  - We will stage new virtual servers for these cases and schedule downtime with campus for migration
- Analysis for hardware requirements and growth are insufficient.
  - 3 Independent sets of analysis have been verified by team
- Migration can be completed in allotted time.
  - Migration from physical to virtual will be ongoing throughout project timeline (currently in progress)

Milestones

- October – November: Assessment of current environment –COMPLETED
- November : Initial hardware specifications – COMPLETED
- December: IDR Presentation – COMPLETED
- January: Bid equipment to vendors
- January: Socialize GWC server/storage project plan with campus stakeholders
  - **Start of Spring Semester 1/31/15**
- February: Procure equipment
- March: Start communication to GWC campus stakeholders of upcoming project plan
- March: Receive and stage equipment
- March: Hardware/Software training
- April: Update GWC campus stakeholders on progress and project plan
- April: ESXi, vCenter, HP OneView install/setup
- May: Update GWC campus stakeholders on progress and upcoming steps (communicate needed downtime of services etc…)
  - **End of Spring Semester 5/31/15**
- June: Migrate and consolidate
  - **2nd week of June session starts**
- August: Summer Session ends
- October: Decommission equipment from MDF

Snapshot of Current Environment

- **EMC cx500 SAN 4gb Fiber Channel (EOL)**
  - Currently holds data for GWC campus shares/email/virtual environment
  - No warranty for hardware or software on production SAN
- **Dell PowerEdge servers (EOL)**
  - Current GWC physical servers
  - Average age is 6 years old
    - Oldest are over 10 years old
- **VDI Environment (Production)**
  - 2 Dell PowerEdge r810
    - Still under warranty from Dell
  - Compellant SAN (4TB useable data)
    - Still under warranty from Dell
  - Cisco Fabric Switches x2
    - Still under warranty from Cisco
  - Approximately 40 Thin/Zero clients in environment
  - Possibly migrate to Blades/3Par in future (2016 +)
- **GWC Backup Infrastructure**
  - Exagrids (x2)
  - Overland Tape Library
- **APC Symmetra Single Phase Battery Backup**
  - Use existing GWC hardware to offset power usage on current District APCs
  - Ability to increase run time between possible power outage and generator start up

GWC Hardware Standardization

- **Standardize the Districts infrastructure in order to focus the Infrastructure team's skill set as well as reduce overall complexity and cost of future District infrastructure**
  - **Provide flexibility, scalability, stability, and resiliency to GWC and Districts Infrastructure**

- **Server Standard**
  - HP C7000 Blade Chassis
    - Current standard for District
  - HP G9 Half Height Blades x8
    - Memory amount not finalized yet
    - Model not yet finalized

- **Storage Standard**
  - HP 3Par SAN 7400
    - Current standard for District
  - Fiber Channel 8gb Connections
    - Roadmap for NFS,CIFs, and Dedupe/Compression in Q1 2015
      - Provides space saving (Dedupe/compression) and protocol flexibility (NFS/CIFs)

- **District APC Racks/Power/Cooling**
  - No additional racks, power, or cooling needed due to District decommissioning plan of hardware in place currently

Network Standardization

- **GWC will leverage existing District infrastructure life cycles and use pre-existing equipment to reduce initial cost of implementation**

- **Fabric Network**
  - Use existing Cisco 9148 fabric switches for 8gb Fiber Channel (FC) connections (GWC servers to GWC storage)
  - Ability to leverage current District infrastructure for GWC fabric network as well as increasing 4x throughput compared to current GWC fabric

- **Ethernet Network**
  - Use existing Dell Force10
  - Leverage layer 2/3 capabilities
  - Possibly use Districts class B address or GWC private address space

District / GWC Network Layout – Option 1

Golden West Network Overview
Flat Network

The Option 1 design would move GWC's servers to the District IP address space. This would require significantly more work but achieves the overall goal of flattening the network.

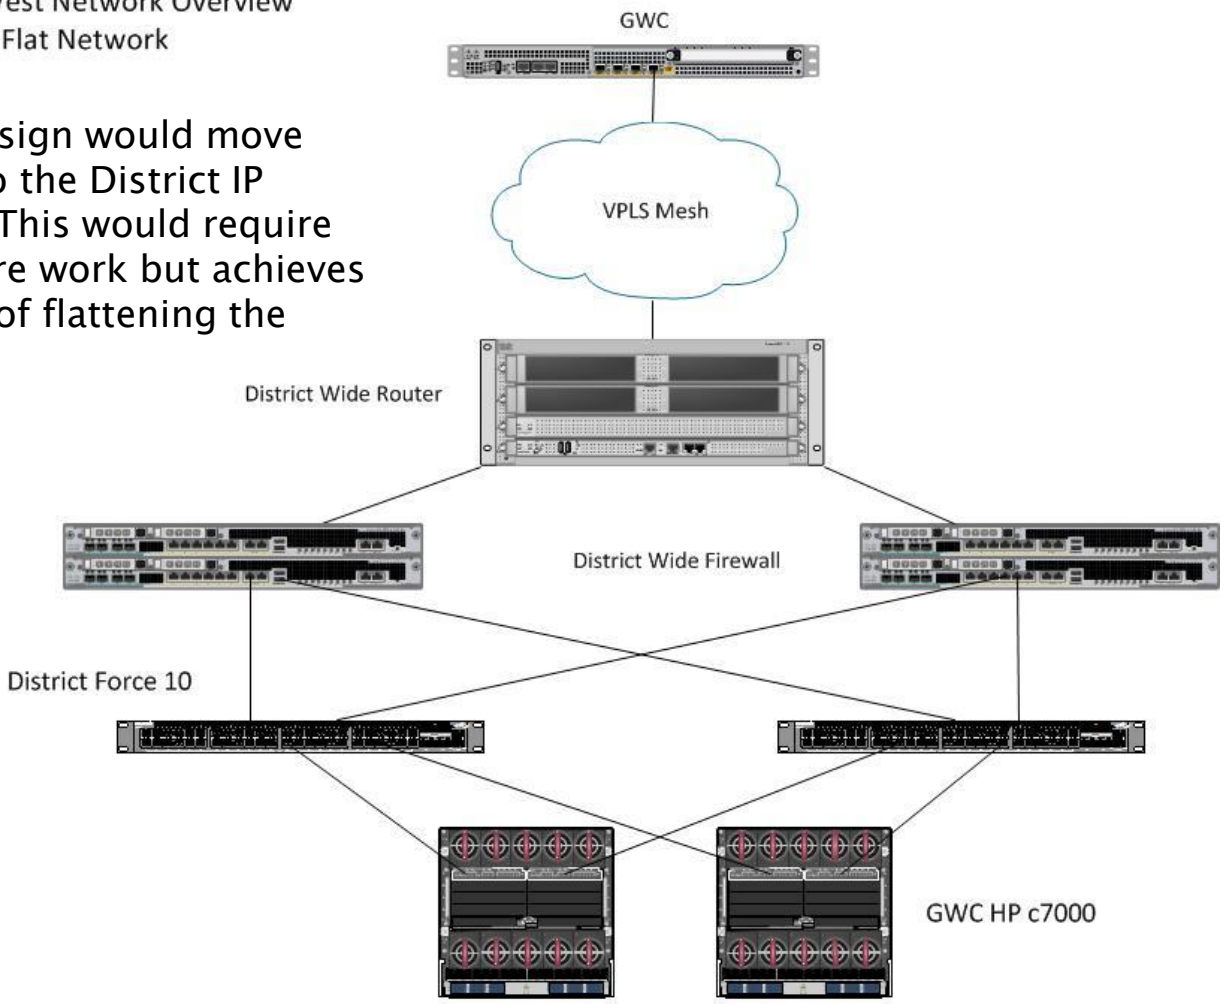

REVISIONS			
REV 1	Golden West Flat Network Overview	DATE 11/14/2014	Updated: KM

District / GWC Network Layout – Option 2

Golden West Network Overview

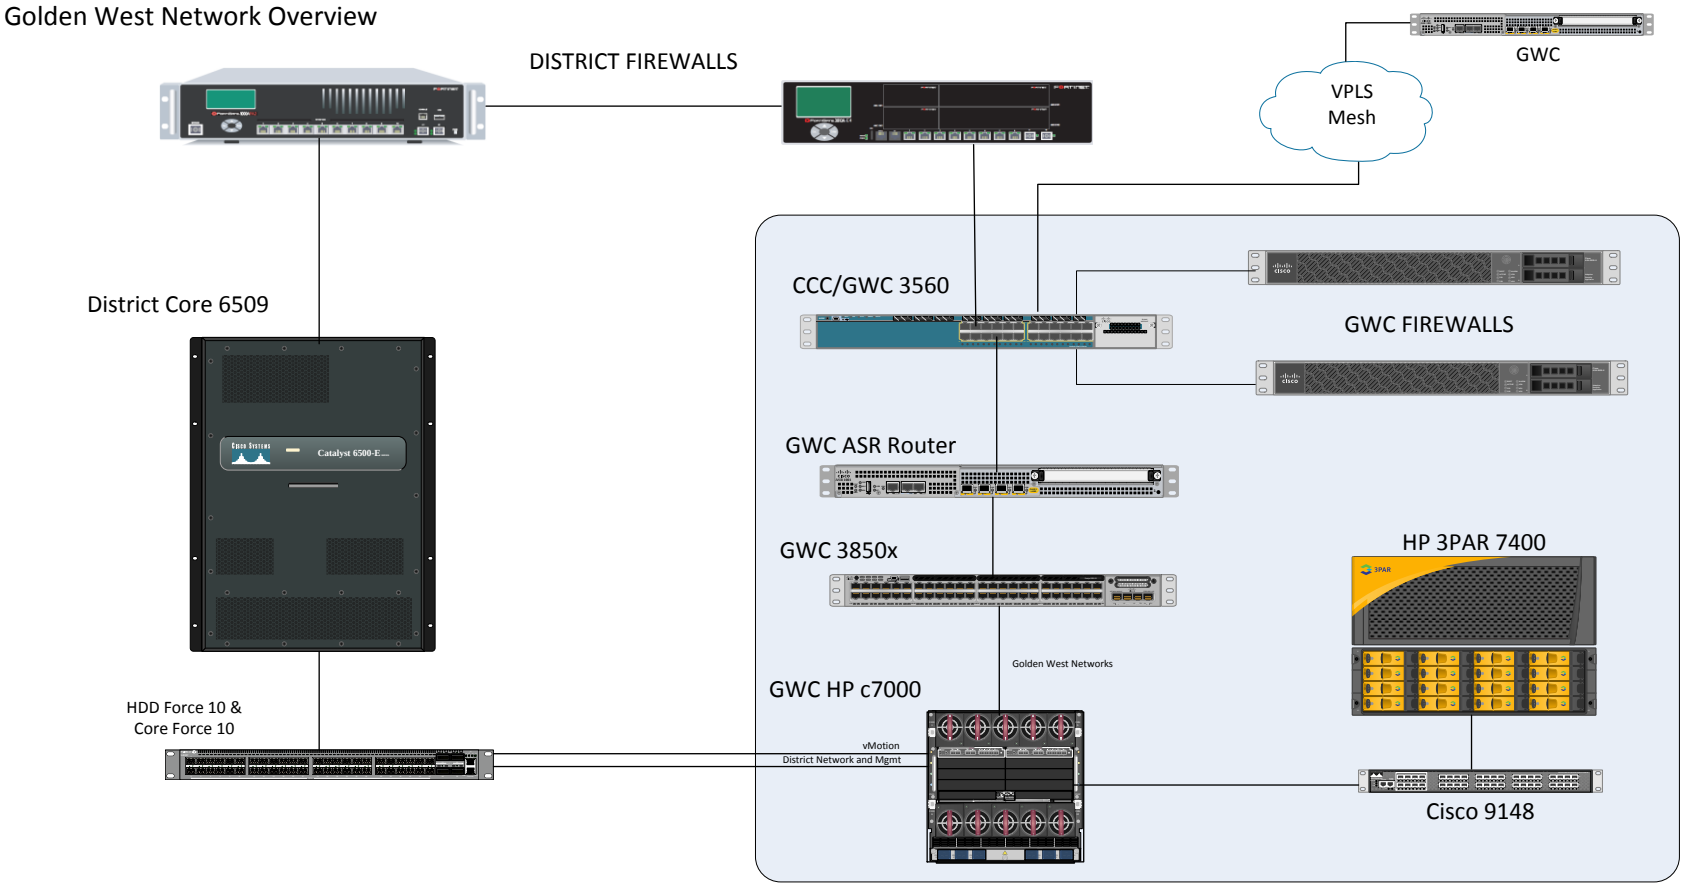

REVISIONS			
REV 2	Golden West Network Overview	DATE 11/20/2014	Updated: KM

The Option 2 design would extend GWC's network into the District DC. IP addressing of the servers would remain the same.

Power and Cooling

- **The capacity of the APC Symmetra UPS at the District will be increased to handle the load from new hardware** (ref slide 25)
  - Use existing GWC MDF Symmetra power modules (2 or 4 of them) and augment current District APC
  - Possibly reuse batteries from current GWC PBX for expanded battery runtime needs at District
    - Can remove batteries during production and will not affect hardware
- **Current 3 x 10 ton AC units will be sufficient for new equipment based on current usage** (ref slide 25)
  - Coastline and OCC consolidation may need to revisit cooling requirements for sufficient AC coverage
- **Racks 13, 16, and 17 have 50 amp 3 phase power**
  - The correct number of and type of outlets are available in these racks for the c7000 and the 3Par
- **VDI and backup equipment will be housed in 25 amp racks**

District Data Center Rack Layout

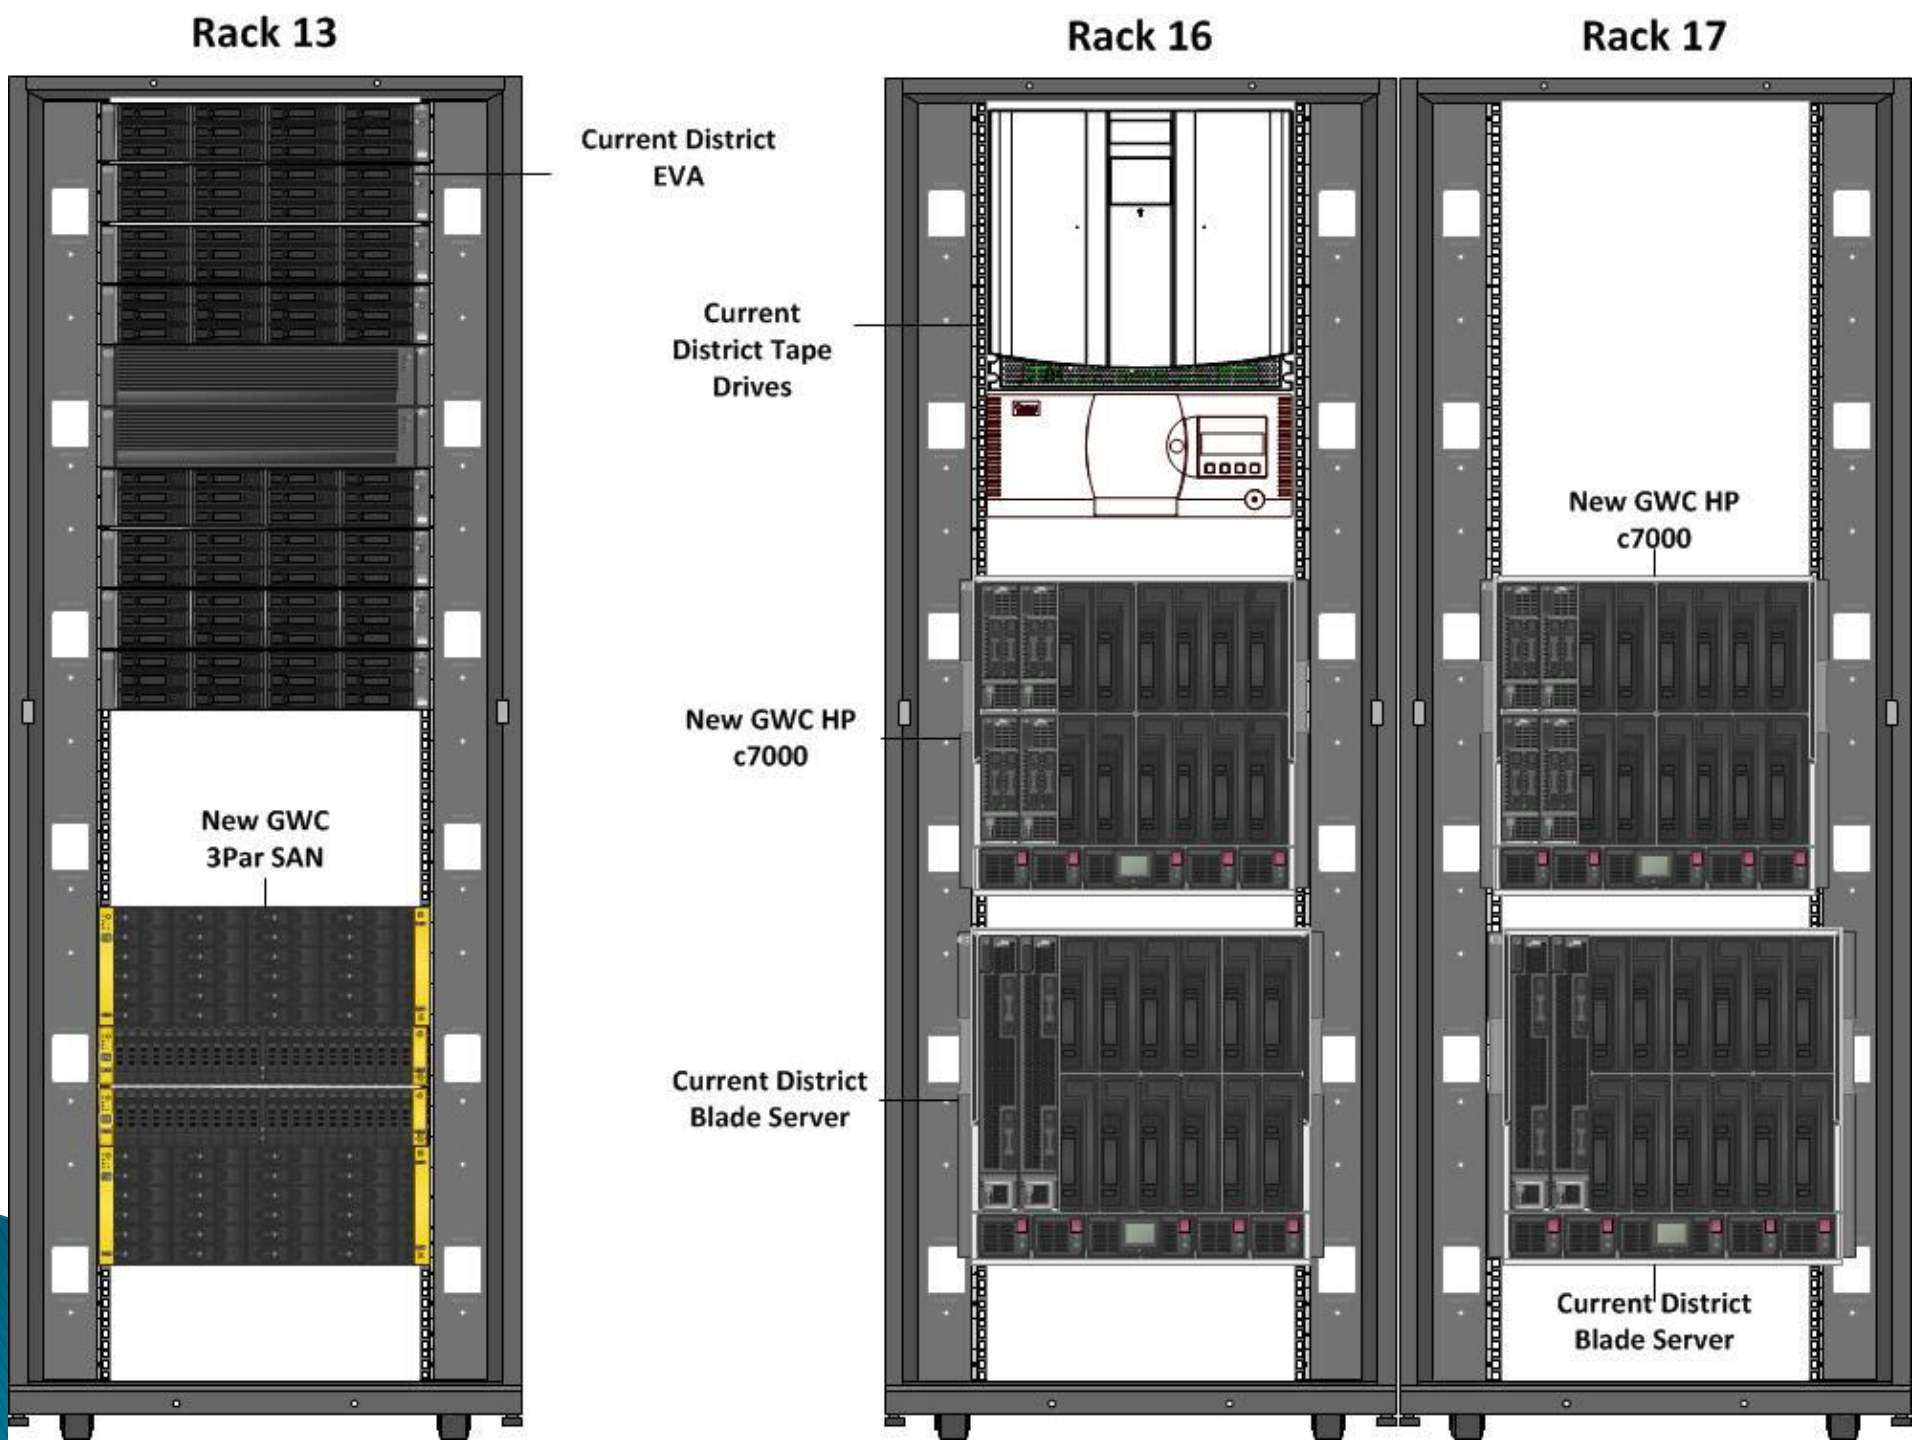
Rack 13
Rack 16
Rack 17
Current District EVA
Current District Tape Drives
New GWC HP c7000
New GWC 3Par SAN
New GWC HP c7000
Current District Blade Server
Current District Blade Server

Redundancy and HA/FT at Virtual Layer

- Physical blade servers will be built for 60% utilization of dedicated resources (memory and processing) of the current virtual machines as well as the potential for growth (approximately 25%). The hypervisor will use 10% of the physical resources for management "overhead" and we build in 30% for peak and burst computing needs. This architecture enables GWC to have multiple host for full N+1 redundancy to accommodate for full chassis and/or blade failure

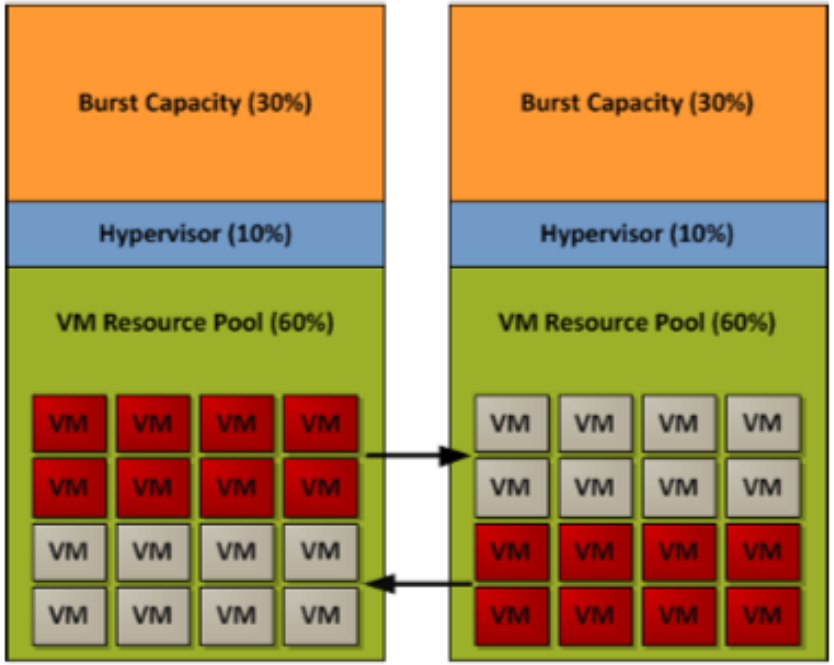

Management of Infrastructure

- **VMware vSphere with Operations Management**
  - Provides performance monitoring
  - Capacity management
  - SPoG management for VMware environment
    - One console to manage all VMware environment
- **HP One View**
  - SPoG management for HP infrastructure District wide
    - One console to manage all HP hardware (Servers/Storage)
    - Firmware updates/management
    - Chassis/Blade template and auto provision based on profiles
    - Administrate and manage HP 3Par SAN, c7000 chassis, and individual blades
  - Plugs in to VMware for seamless management and reporting
- **WUG Will be used for network monitoring of District Infrastructure**

Vendor for Hardware/Services

- Vendor A has provided assessment of GWC infrastructure and initial pricing (WAG Pricing)
  - Server (physical and virtual) current metrics
  - Storage metrics (Size, IOPS, Latency)
  - Vendor A has made recommendations based on our metrics and industry best practice
- District IT and Vendor B have provided metrics as well for review and validation of hardware needs for current and growth (estimated growth of 25% over 3 years) (ref slide 22, 23, and 24)
  - Assumption based on addition of 4k-5k new students (~1000 FTE) by EOY 2015
  - Based on decommissioning of GWC Exchange, ADM domain, INS domain, student file shares

HP c7000 Chassis and Gen9 Blades

	Qty.	Part	Description	
Hardware	1	681844-B21	HP BLC7000 CTO 3 IN LCD PLAT ENCLOSURE	
	8	727021-B21	HP BL460C GEN9 10GB/20GB FLB CTO BLADE	
	8	726990-L21	HP BL460C GEN9 E5-2660V3 FIO KIT	
	8	726990-B21	HP BL460C GEN9 E5-2660V3 KIT	
	128	726719-B21	HP 16GB 2RX4 PC4-2133P-R KIT	
	8	700764-B21	HP FLEXFABRIC 20GB 2P 650FLB FIO ADPTR	
	8	761871-B21	HP SMART ARRAY P244BR/1G FIO CONTROLLER	
	2	571956-B21	HP BLC VC FLEXFABRIC 10GB/24-PORT OPT	
	1	517521-B22	HP 6X 2400W PLAT HT PLG FIO PWR SPLY KIT	
	1	456204-B21	HP BLC7000 DDR2 ENCL MGMT OPTION	
	1	677595-B21	HP BLC 1PH INTELLIGENT POWER MOD FIO OPT	
	1	517520-B21	HP BLC 6X ACTIVE COOL 200 FIO FAN OPT	
				Extended Sell
			Hardware Total:	**$128,296.46**

HP c7000 and 8 Gen9 Blades
~ $130,000

HP 3Par, 3 tier, 37+TB usable

	Qty.	Part		Description	
Hardware	1	QR485A		HP 3PAR STORESERV 7400 4-N STORAGE BASE	
	4	QR486A		HP 3PAR 7000 4-PT 8GB/S FC ADAPTER	
	24	QR494A		HP M6710 450GB 6G SAS 10K 2.5IN HDD	
	16	E7W54B		HP M6710 480GB 6G SAS 2.5IN MLC 5YR SSD	
	2	QR490A		HP M6710 2.5IN 2U SAS DRIVE ENCLOSURE	
	2	QR491A		HP M6720 3.5IN 4U SAS DRIVE ENCLOSURE	
	24	QR499A		HP M6720 2TB 6G SAS 7.2K 3.5IN NL HDD	
	1	QR516B		HP 3PAR 7000 SERVICE PROCESSOR	
	1	BD362AAE		HP 3PAR STORESERV MGMT/CORE SW E-MEDIA	
	1	BD363AAE		HP 3PAR 7000/7450 OS SUITE E-MEDIA	
	1	BD365AAE		HP 3PAR 7000 SERVICE PROC SW E-MEDIA	
	1	BD368AAE		HP 3PAR APP SUITE FOR EXCHANGE E-MEDIA	
	1	BD371AAE		HP 3PAR APP SUITE FOR SQL E-MEDIA	
	1	BD372AAE		HP 3PAR APP SUITE FOR VMWARE E-MEDIA	
	1	BD373AAE		HP 3PAR REPORTING SUITE E-MEDIA	
	8	QK734A		HP PREMIER FLEX LC/LC OM4 2F 5M CBL	
					Extended Sell
				Hardware Total:	**$117,380.00**
Software	1	BC795B		HP 3PAR 7400 REPORTING SUITE LTU	
	1	BC796B		HP 3PAR 7400 APP SUITE VMWARE LTU	
	1	BC797B		HP 3PAR 7400 APP SUITE EXCHANGE LTU	
	1	BC798B		HP 3PAR 7400 APP SUITE SQL LTU	
	1	BC773B		HP 3PAR 7400 OS SUITE BASE LTU	
	64	BC774A		HP 3PAR 7400 OS SUITE DRIVE LTU	
	1	BC775A		HP 3PAR 7400 REPLICATION SUITE BASE LTU	
	64	BC776A		HP 3PAR 7400 REPLICATION SUITE DRIVE LTU	
	1	BD270A		HP 3PAR 7400 DATA OPT ST V2 BASE LTU	
	64	BD271A		HP 3PAR 7400 DATA OPT ST V2 DRIVE LTU	
					Extended Sell
				Software Total:	**$47,728.01**
Support	1	H1K92A3	RDD	HP 3PAR 7400 OS SUITE BASE LTU SUPP	
	1	H1K92A3	RDE	HP 3PAR 7400 RPLICATIONSUITE BSE LTU SUP	
	1	H1K92A3	RDQ	HP 3PAR 7400 REPORTING SUITE LTU SUPP	
	3	H1K92A3	RDR	HP 3PAR 7400 APP SUITE LTU SUPP	
	1	H1K92A3	RZ5	HP 3PAR 7000 SERVICE PROCESSOR SUPP	
	64	H1K92A3	S7D	HP 3PAR 7400 OS SUITE DRIVE LTU SUPP	
	64	H1K92A3	S7E	HP 3PAR 7400 REPLIC SUITE DRIVE LTU SUPP	
	1	H1K92A3	S7W	HP 3PAR 7400 DATAOPT ST V2 BASE LTU SUPP	
	64	H1K92A3	S7X	HP 3PAR 7400 DATA OPT ST V2 DRV LTU SUPP	
	16	H1K92A3	TGD	HP 3PAR7000 480GB SAS SSD SUPP	
	24	H1K92A3	WUS	HP 3PAR 7000 DRIVES UNDER 1TB SUPPORT	
	24	H1K92A3	WUT	HP 3PAR 7000 DRIVES OVER 1TB SUPPORT	
	4	H1K92A3	WUW	HP 3PAR 7000 DRIVE ENCLOSURE SUPPORT	
	4	H1K92A3	WUX	HP 3PAR 7000 ADAPTER SUPPORT	
	1	H1K92A3	WUY	HP 3PAR 7400 4-NODE STORAGE BASE SUPPORT	
					Extended Sell
				Support Total:	**$53,673.67**
Services	1	HA114A1	5TS	HP STARTUP 3PAR 7400 4-ND STRG BASE SVC	
	4	HA114A1	5TT	HP STARTUP 3PAR 7000 FC ADAPTER SVC	
	2	HA114A1	5TV	HP STARTUP 3PAR 7000 2U SAS ENCLOSRE SVC	
	2	HA114A1	5TW	HP STARTUP 3PAR 7000 4U SAS ENCLOSRE SVC	
					Extended Sell
				Services Total:	**$5,655.36**
					Extended Sell
				Solution Total:	**$224,437.04**
				Grand Total:	**$224,437.04**

HP 3PAR 7400

~ $230,000

GWC Server and Storage Consolidation Cost

Description	Cost
HP c7000 Chassis + Gen 9 blades	~$130,000
HP 3Par 7400	~$230,000
Miscellaneous (patch cables, SFP, Professional Services, etc)	~$35,000
Total	~$400,000

Reference Slides

Golden West Assessment Results

- Assessments shows average CPU and memory usage averaging 4.09% and 47% respectfully (ref slide 23, 24, and 25)
- Assessments shows average IOPS Read/Write aggregate of 1425 (ref slide 23 and 24)
- Approximate 40% read IOPS and 60% write IOPS
  - Burst IOPS peaks 3707 per second (ref slide 23)

Golden West Server Utilization

Workload Estimates	Value
Average CPU count	4
Average CPU per servers (MHz)	2,710
Average CPU per server 4CPU x 2,710MHz	10,840
Average CPU utilization (%)	4.09%
Workload average peak CPU utilization (MHz) 10,992 x 4.09%	443
Total Peak average CPU utilization for a 80 workload (MHz)	35,468
Average RAM per server (MB)	5,012
Avarage RAM utilization (%)	47.00%
Workload average peak RAM utilization (MB) 5,012 MB x 47%	2,356
Total Peak RAM Utilization for a 80 workload (MB)	188,451

The current virtual and physical infrastructure at Golden West requires 35 Ghz of CPU and 188 GB of RAM at average peak utilization.

To account for 25% of growth the server cluster should be sized for 44 Ghz of CPU and 235 GB RAM.

NOTE: CPU and RAM count was calculated from manual inventory of virtual and physical servers. Utilization numbers were acquired from MAP report.

Golden West Vendor B DPACK Results

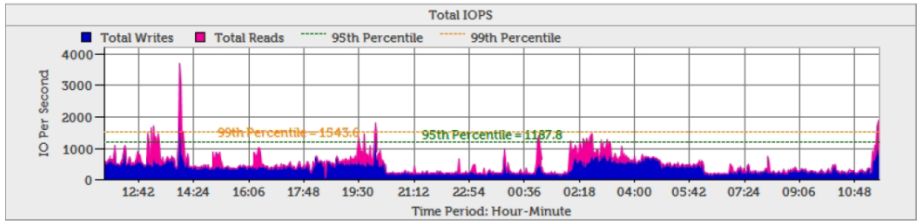

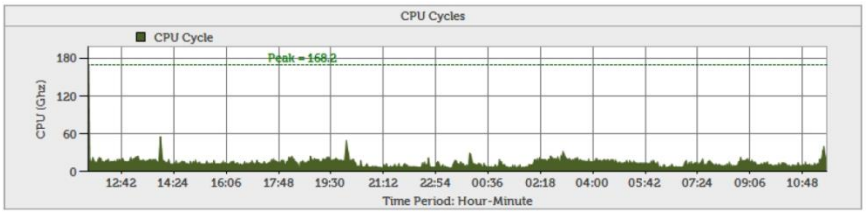

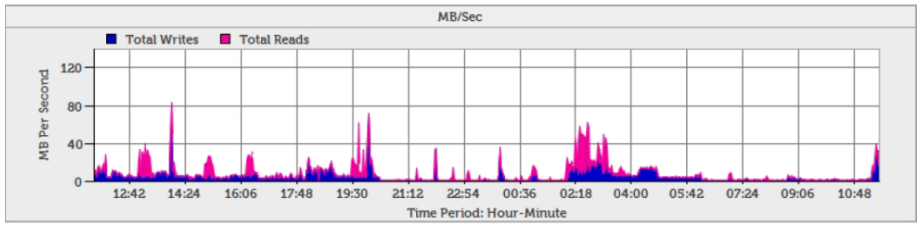

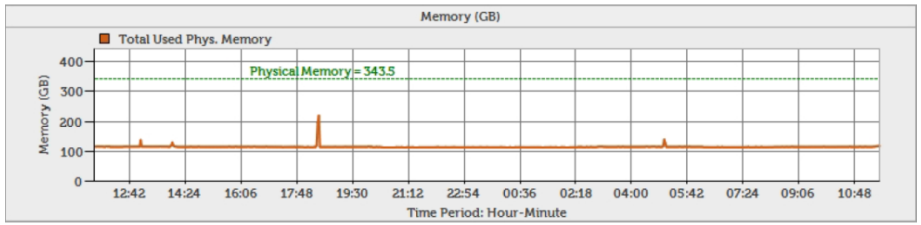

DPACK is a tool that captures the IOPS, CPU, and Memory utilization of a pool of servers. The aggregation of these metrics can then be used to determine peak utilization.

Vendor B's DPACK report illustrates a total of 344 GB RAM and 168.2 Ghz of CPU for 69 workloads.

It also illustrates RAM utilization on average is well below 200 GB and CPU utilization is well below the 60 GHz line.

Aggregate Output	
Time Recorded	23 Hour(s) 56 Minute(s), 9/23/2014 - 9/24/2014
Servers Recorded	69
Output Summary:	
Throughput (MB/s)	**83.3**
IOPS	**1187.8 at 95%, 1543.6 at 99% and 3707.4 at peak**
Read/Write Ratio	42% and 58%
Total Capacity:	19981.5 GBs
Free Capacity	8114.8 GBs Free
Used Capacity	11866.6 GBs Used
Total Cores	259
Total Processors	211
CPU Aggregate	168.2 Ghz

Vendor A MAP Assessment

Environment Information	
CPU	
Peak CPU Utilization	15.20%
Average CPU Utilization	4.09%
Memory	
Total Available (GB)	321
Total Peak Usage (GB)	147
Overall Usage	46%
Average Memory used per System	47%
Systems over 90% Memory Usage	1

Storage Sizing	
	All Servers
Physical Drive Total (GB)	12,324
Logical Drive Used (GB)	10,012
95th Percentile	
95th Percentile Reads per Sec	583
95th Percentile Writes per Sec	2977
95th Percentile Transfers per Sec	3540
Minimum Configuration	
Logical Drive Used w/ Growth (GB)	12,015
Storage Growth	20%
RAID 5 w/ 15k Drives	70 drives at minimum size of 174G each
Total Average Read IOPS	583
Total Average Write IOPS	842

Hardware Recommendations	
Sizing Information	
Total MHz (Mix of Peak and Average)	807.16
Total Peak Memory Used (GB)	147
Network Observations	
Peak Network Traffic Total (Mb/sec)	1
Highest Network Traffic System (Mb/sec)	0
Ports per Host (minimum 1 Gb throughput)	
Network Ports for Virtual Machines	2
Additional Ports for best practice	6
Server Hardware Proposed #1	
Proposed System	HP DL380p Gen8 (E5-2600 v2)
Proposed Quantity	3
Target Cluster Utilization	60%
Growth Potential	100%
CPU	
CPU Speed(MHz)	1800
Number of Sockets	2
Cores per Socket	8
Memory	
RAM per system (GB)	256

MAP assessment is Microsoft Assessment Planning toolkit. MAP securely assess IT environment for platform migration. MAP gathers metrics based on IOPS, CPU, and Memory utilization over a period of time (30 days in GWC instance).

Power and Cooling Reference

- Approximate GWC power consumption will increase the District's heat load by 2.88 kW
  - Currently District creates ~63 kW
  - Currently GWC physical environment uses approximately 8.38 kW
    - Consolidation reduces consumption almost 3 times
- District uses approximately 18 of the 30 ton capacity
  - Total Data Center CRAC capacity is 3 x 10 tons
  - GWC will increase HVAC consumption to 18.8 tons
    - Formula used 3.5kW = 1 Ton of HVAC (rough estimate)
  - District HVAC is currently over subscribed for N+1 redundancy
    - Will have to be discussed for future of District DC expansion and campus consolidation efforts