The University of Utah
Data Center Conceptual Master Planning Project

Final Report

July 31, 2009

Project Ref: 5K3-UU001

HP Critical Facilities Services

delivered by EYP MCF

Document information

Title, Subject	Unified Conceptual Master Plan
Author:	Steve Carter, Rob Myers
Release Date	31 July 2009
Version	1
Prepared for	University of Utah
Document Name	U_of_U_Conceptual_Master_Plan_v1_Issued.pdf
Classification	

Change History

Version	Date	Comments	Initials

Approvals

Name	Signature	Title	Date

Signed acceptance forms are filed in the EYP MCF project files.

TABLE OF CONTENTS

1 EXECUTIVE SUMMARY 5

1.1 OVERVIEW 5
1.2 FUTURE STATE PLANNING CONCEPTS 6
1.3 PROJECT COLLABORATION 7
1.4 PROGRAM RECOMMENDATION 9

2 ENTERPRISE CORE COMPUTING CONCEPTS 10

2.1 OVERVIEW 10
2.2 KNOWN CAPACITY ISSUES 11
2.3 MARRIOTT LIBRARY CAPACITY GAP ANALYSIS 12
2.4 PHYSICAL DATA CENTER PLANNING – ENTERPRISE CORE 13

3 CENTER FOR HIGH PERFORMANCE COMPUTING CONCEPTS 16

3.1 OVERVIEW 16
3.2 PHYSICAL DATA CENTER PLANNING – HPC 18

4 CO-LOCATION 19

4.1 OVERVIEW 19
4.2 PHYSICAL DATA CENTER PLANNING – CO-LOCATION 20

5 DATA CENTER PROGRAM RECOMMENDATION 21

5.1 OVERVIEW 21
5.2 OPPORTUNITIES AND CHALLENGES - PROGRAM RECOMMENDATION 22
5.3 PHYSICAL DATA CENTER PLANNING – PROGRAM RECOMMENDATION 23
5.4 PROGRAM RECOMMENDATION – TEST FIT 24
5.5 MECHANICAL AND ELECTRICAL CONCEPTS 25
5.6 PROGRAM RECOMMENDATION ROM COST ESTIMATE 26

6 BEST PRACTICES 27

6.1 CRITICAL ISSUES BEST PRACTICES GAP ANALYSIS 27
Disaster Recovery / Business Continuity Planning 27
Application Criticality Prioritizing Organized by Business Unit Requirements 28
Data Replication and Storage Standards 28
High Growth Area – High Performance Computing (HPC) 29
IT Systems Deployment Standards 30
6.2 IT OPERATIONS BEST PRACTICES RECOMMENDATIONS 31
6.3 LIMIT PHYSICAL ACCESS TO THE PRODUCTION DATA CENTER 31
6.4 CONSOLIDATION OF MULTIPLE SMALLER DATA CENTERS INTO LARGER DATA CENTER 32
6.5 RECOMMENDED NEXT STEPS 34

7 DATA CENTER MASTER PLANNING METHODOLOGY DISCUSSION 35

7.1 FUTURE STATE SPACE, POWER AND COOLING REQUIREMENTS 35
Ten Year Space, Power and Cooling Requirement Estimates and Assumptions 35
Space Planning Terms and Definitions 36
7.2 DATA CENTER PLANNING CONSIDERATIONS 37
Overview 37
Standardization 37
Deployment Model 37
7.3 BUSINESS CONTINUITY PLANNING AND DISASTER RECOVERY CONSIDERATIONS 47
BCP Overview 47
Application Prioritization Ratings 49
Attributes of enhanced BCP/HA/DR capabilities utilized in this study include: 49
Network Based Operational Continuity Considerations 50
7.4 DATA CENTER FACILITIES CONSIDERATIONS 51
Staging Area 51
Storage 52
Physical Security 52
Loss Prevention/Risk Management 52
Fire Protection 53
Piping and Drainage 53
Environmental Issues 53
Acoustical Considerations 53
Employee Welfare Issues 53
Flexibility of Use 54
LEED 54
7.5 THE UPTIME INSTITUTE CLASSIFICATION (FROM THE UPTIME INSTITUTE) 55
Tier I - Basic Non-Redundant Data Center 55
Tier II - Basic Redundant Data Center 55
Tier III - Concurrently Maintainable Data Center 56
Tier IV - Fault Tolerant Data Center 56
7.6 DATA CENTER CONCURRENTLY MAINTAINABLE REDUNDANCY ATTRIBUTES 56
7.7 DATA CENTER MAINTENANCE ASSURANCE 57
7.8 POWER USAGE EFFECTIVENESS (PUE) AND DATA CENTER INFRASTRUCTURE EFFECTIVENESS (DCIE) 59
PUE Estimates for Data Center Facilities 62
7.9 HIGH AVAILABILITY STORAGE REPLICATION 64
Recovery Time Objective 64
7.10 HIGH DENSITY COMPUTING 65
7.11 DATA CENTER MIGRATION STRATEGY 66
Approach 66
Conclusion 68

8 APPENDIX A – GLOSSARY 69

9 APPENDIX B – DATA CENTER WORKSHOP PLANNING NOTES 71

1 Executive Summary

In May of 2009, University of Utah engaged the Hewlett-Packard Company to create a data center conceptual master plan. The primary objective of the project was to evaluate the current state of the seven existing data centers identified by University of Utah and to provide recommendations for sizing of a data center to accommodate consolidation of these data centers into an existing building located at 875 S. West Temple Street, Salt Lake City. The work was performed by EYP Mission Critical Facilities, Inc. (EYP MCF), a subsidiary of HP.

The seven data centers identified by University of Utah were:

- Park Building
- Student Services Building
- Eccles Broadcast Center
- Fort Douglas
- Komas Datacenter
- Komas Cluster
- Marriott Library

1.1 Overview

EYP MCF performed our data center planning workshops which focus on how business requirements and Information Technology current and future state planning would impact design requirements for the new data center. These workshops included consideration of the data processing hardware and the physical space, power and cooling requirements of existing and planned data center facilities. The objectives included five-year space, power, cooling requirements for University of Utah's data center infrastructures, and the consolidation of existing data centers. EYP MCF did NOT perform mechanical and/or electrical assessments of the existing data centers and relied solely on stakeholder input. This project focused on physical space, power and cooling requirements and high-level data center conceptual future state recommendations, taking into account the University of Utah Data Center Improvements Programming requirements.

EYP MCF utilized input from stakeholders/user groups collected during these workshops and internal University of Utah reports to develop forward-looking space, power and cooling requirements. This essential input from University of Utah stakeholders was utilized in the development of all concepts and recommendations provided within this report. During the interviews, EYP MCF discussed the core topics of Disaster Recovery/Business Continuity, current infrastructure status, future infrastructure needs, current IT state, and future IT state with University of Utah stakeholders.

Analysis and modeling was based on conclusions resulting from stakeholder meetings.

- 24/7 access required by University of Utah customers causing extreme difficulty for scheduling maintenance.
- Would need substantial bandwidth from campus and Richfield to new data center to accomplish intended service offerings, disaster recovery and high availability.
- EYP MCF agrees with the University's conclusion that the current data center topology will not meet the University's future needs and a purpose built data center is needed.

- High performance computing systems have unique requirements compared to University core systems and is best served in a separate environment designed to different reliability standards.

The following concepts were evaluated which led to the Program Recommendation by EYP MCF and University of Utah stakeholders.

1.2 Future State Planning Concepts

University of Utah faces serious challenges to its ability to meet future data and computing needs. In order for the University to meet the demands for high performance computing, maintain business continuity, and remain competitive for research grants among other universities, University of Utah should seriously consider the recommendations put forward by EYP MCF in this Master Plan.

The major recommendations in this report are based on analysis of the following considerations:

- Enterprise Core Computing
  - U of U is currently experience significant difficulty supporting the space, power and cooling requirements of new Enterprise specific systems within the current datacenter infrastructures.
  - Current growth projections indicate current data centers cannot accommodate anticipated needs.
  - Any future data centers should be concurrently maintainable due to the increasing demand on nearly all applications requiring 24/7 access.
  - Optimization efforts continue to be implemented and utilized whenever possible.
  - Facility should operate with dedicated operations staff with limited access to data center floor.
  - Build out should be developed in phases.
  - Richfield will continue to be disaster recovery site
  - Existing ACS and Marriott Library data centers may be utilized in HA scenarios.

- Center for High Performance Computing
  - University of Utah is currently experiencing significant difficulty supporting the space, power and cooling requirements of CHPC specific systems within the current datacenter infrastructures.
  - Significant growth is anticipated for CHPC specific systems over the next five years. Growth of 100% is common for CHPC systems every 2-3 years within large R1 Universities.
  - High performance computing systems have unique requirements compared to University core systems and is best served in a separate environment designed to different reliability standards
  - No known solution/plan exists to accommodate long-term growth of CHPC specific systems.

- Removal of HPC type systems from the existing datacenters could provide some interim relief from existing datacenters current capacity limitation issues.
- Consensus recommendation is to design and build a separate CHPC specific datacenter space.

- Collocation
  - Co-Location offering is a concept the University is interested in offering colleges and other entities with University relationships.
  - Possibility of new Fiber Optic MAN opens the possibility for a Co-Lo offering.
  - Specific needs have not yet been quantified.
  - High-level analysis of co-location offering has deemed the offering **not viable for initial build out** of new data center.

1.3 Project Collaboration

EYP MCF gathered input across the University from numerous knowledge owners, stakeholders, and user groups as well as University documents to produce this report. Groups from the University of Utah organization that participated in the stakeholder interview sessions:

- IT Organization – All major areas organized by technology delivery types:
  - ITS
  - OIT
  - ACS
  - UEN
- Facility Operations
- Executives
- Center for High Performance Computing
- Applications Development Leadership:
  - Clinical Systems
  - Enterprise Systems
  - ACS Systems

During the interviews, EYP MCF discussed the following core topics with University of Utah provided stakeholders:

- Current IT architecture and infrastructure operational and planning practices.
- Disaster recovery and/or Business Continuity plans.
- Data replication and storage standards.

- IT systems deployment standards.
- IT systems growth trends and future state analysis.
- Vision of University of Utah leaders for the future state of IT systems.
- Business growth projections with associated IT systems growth requirements.
- Consolidation of University of Utah production data centers.
- Network architecture/infrastructure initiatives and future state plans.
- Utilization of virtual server platforms to attempt to manage significant growth of distributed computing platforms.
- Operational capabilities of the University of Utah IT organization.
- High Performance Computing (HPC) and its effect on existing data centers.

EYP MCF found that University of Utah stakeholders were willing to discuss these issues in detail. Both initial and follow up meetings were held with University of Utah stakeholders to discuss these important topics. The body of this report discusses in detail the concepts and recommendation that were developed through collaboration between EYP MCF and University of Utah primary stakeholders. This collaborative effort was important input for the development of the Program Recommendation detailed in this report.

1.4 Program Recommendation

It is EYP MCF's recommendation that the University of Utah develop and implement a purpose built data center to meet future needs. Of the seven current data centers identified by the University, only Marriott Library has some limited available capacity. All others have reached capacity in one or more of the three areas of space, power or cooling; and the Marriott Library only has capacity to accommodate projected growth until mid 2010.

The Program Recommendation is based on consolidating all of the seven data centers into one primary data center with the possibility of utilizing Marriott Library and/or Park Building for applications that require high availability (failover). Richfield has been identified to remain as the Disaster Recovery location and may be utilized for high availability.

The EYP MCF recommendation is to build out the existing facility shell located at 875 S. West Temple Street as a purpose built data center, utilizing the south bay of the existing building, to accommodate two separate areas - Enterprise Core Computing and High Performance Computing (HPC). HPC has unique requirements that cannot be accommodated efficiently in current data center environments. Separation of these systems allows for independent management of these separate spaces.

Enterprise Core Computing Concept:

- Accommodates Enterprise Core Program Build needs and removes the corresponding loads from existing datacenter spaces.
- Richfield will continue to be disaster recovery site.
- Richfield may provide the ability to house some High Availability (HA) systems. HA systems would be split between two data centers.
- Existing ACS and Marriott Library data centers may be utilized for HA scenarios.

HPC Computing Concept:

- HPC unique requirements that cannot be accommodated efficiently in current data center environments.
- Accommodates CHPC Program Build needs and removes the corresponding loads from existing datacenter spaces.
- HPC Core area will utilize Enterprise Core infrastructure.

2 Enterprise Core Computing Concepts

2.1 Overview

During stakeholder meetings, EYP MCF found that the University of Utah is currently experiencing significant difficulty supporting the space, power and cooling requirements of new enterprise specific systems within the current datacenter infrastructure. Current data centers have known infrastructure capacity issues that have caused outages in the past; and capital investment would be required to increase data centers capacity to meet increasing demand on nearly all applications to maintain 24/7 uptime expectations. Considering all of the data centers would require upgrades or complete renovations to meet University requirements, it is logical to move to a purpose built data center to consolidate all of the seven data centers into one.

The current cooperation between the Hospital and Enterprise IT Organizations surpasses other Universities EYP has assessed to date. This cooperation is enabling the University of Utah to take advantage of a shared services model and allows for efficient use of data center resources and above average consolidation. This collaboration allows for the Enterprise Core to be operated and managed as one cohesive space.

With the lack of available capacity within existing data centers, optimization efforts undertaken by the University have enabled the University to meet growing IT load demands. **This condition is not sustainable.** EYP MCF believes the University of Utah is at risk of increased frequency of data center outages due to current facility infrastructure capacity/reliability issues and will be unable to meet continuing data center resource needs in the near future.

The "Enterprise Core" of the program recommendation in this report is based primarily on the following considerations developed jointly by the University and EYP MCF.

- Any future data centers should be concurrently maintainable due to the increasing demand on nearly all applications requiring 24/7 access.
- Optimization efforts will continue to be implemented and utilized whenever possible.
- Facility should operate with dedicated operations staff with limited access to data center floor.
- Build out should be developed in phases – initial program to support at least 5 year growth projections.
- Richfield will continue to be disaster recovery site
- Existing Park Building and Marriott Library data centers may be utilized in HA scenarios.
- UEN would only move select systems to new data center.
- Printing services would not be housed in the new data center.

2.2 Known Capacity Issues

All data centers, with the exception of Marriott Library, are at full capacity in space, power and/or cooling. Current data centers cannot sustain the Universities anticipated growth needs.

Facility	Hosted Departments	Space	Power	Cooling
Park Building	ACS	At Capacity	At Capacity	Limited Capacity Available
Student Services Building	CHPC and OIT	Capacity Available	At Capacity	Limited Capacity Available
Eccles Broadcast Center	OIT	At Capacity	At Capacity	At Capacity
Fort Douglas	OIT and Media Solutions	Capacity Available	Capacity Available	At Capacity
Komas Datacenter	OIT and Health Sciences	At Capacity	Capacity Available	Capacity Available
Komas Cluster	CHPC, OIT and Health Sciences	Capacity Available	At Capacity	At Capacity
Marriott Library	Library, CHPC and OIT	Capacity Available	Capacity Available	Capacity Available

Figure 1: University of Utah Data Centers Known Capacity Issues

2.3 Marriott Library Capacity Gap Analysis

Marriott Library was initially considered to be the "stop gap" data center to sustain growth until the new data center is in production. Current projections show Marriott library will reach capacity during year 2010. Continued diligence as to additional systems added to Marriott Library is necessary to ensure availability for critical applications.

Marriott Library			
	Space	Power	Cooling
Total Build	2700 sq. ft.	120 kW	40 Tons
Current Utilization	225 sq. ft.	12 kW	3.4 Tons
Capacity Remaining	2475 sq. ft.	108 kW	36.6 Tons
Capacity Reached	Early 2013	**Mid 2010**	Early 2011

- Assumptions:
  - Marriott Library will accept all future Enterprise Core growth.
  - Growth projections are based on continued optimization efforts.

Projections indicate Marriott Library will reach capacity and be unable to accept new equipment around the middle of year 2010.

Figure 2: Marriott Library Capacity Gap Analysis

2.4 Physical Data Center Planning – Enterprise Core

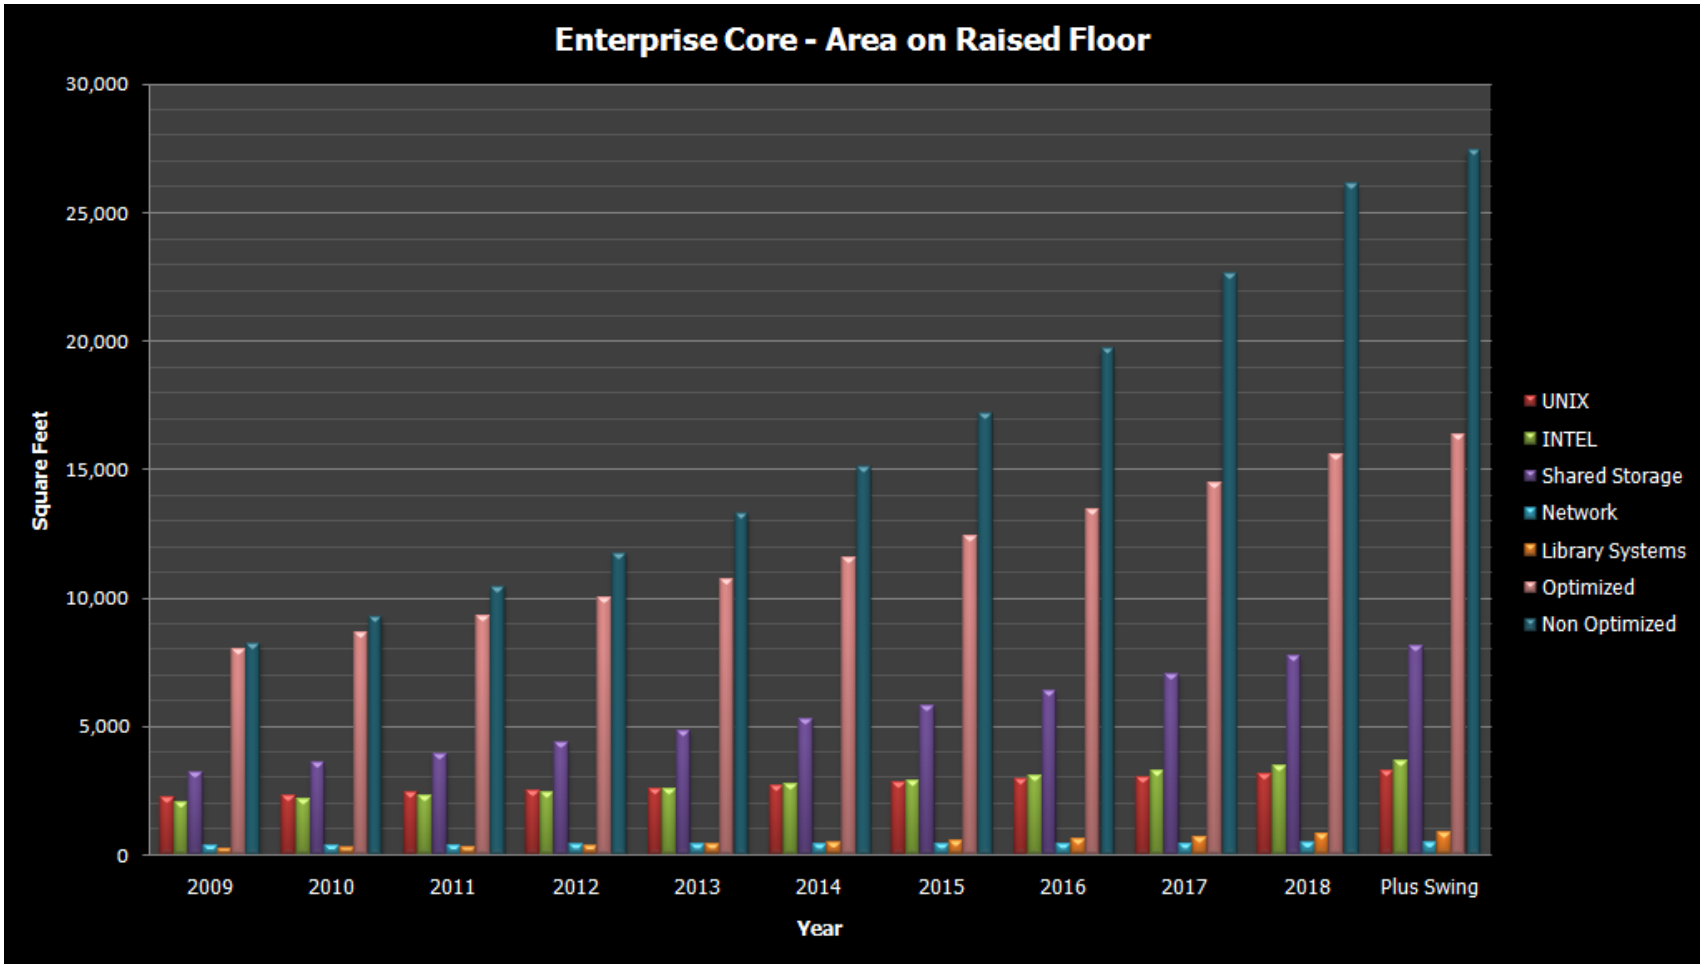

Figure 3: Ten Year Enterprise Core Raised Floor Projections

- The figure above shows that approximately 11,000 sq. ft. of raised floor area would be required for the Enterprise Core 5 year (Day 1) growth planning.
- The figure above shows that approximately 16,000 sq. ft. of raised floor area would be required for the Enterprise Core 10 year (Final) growth planning.
- The optimized model assumes storage technologies will continue to improve data storage to raised floor footprint ratios and the University applies these technologies.

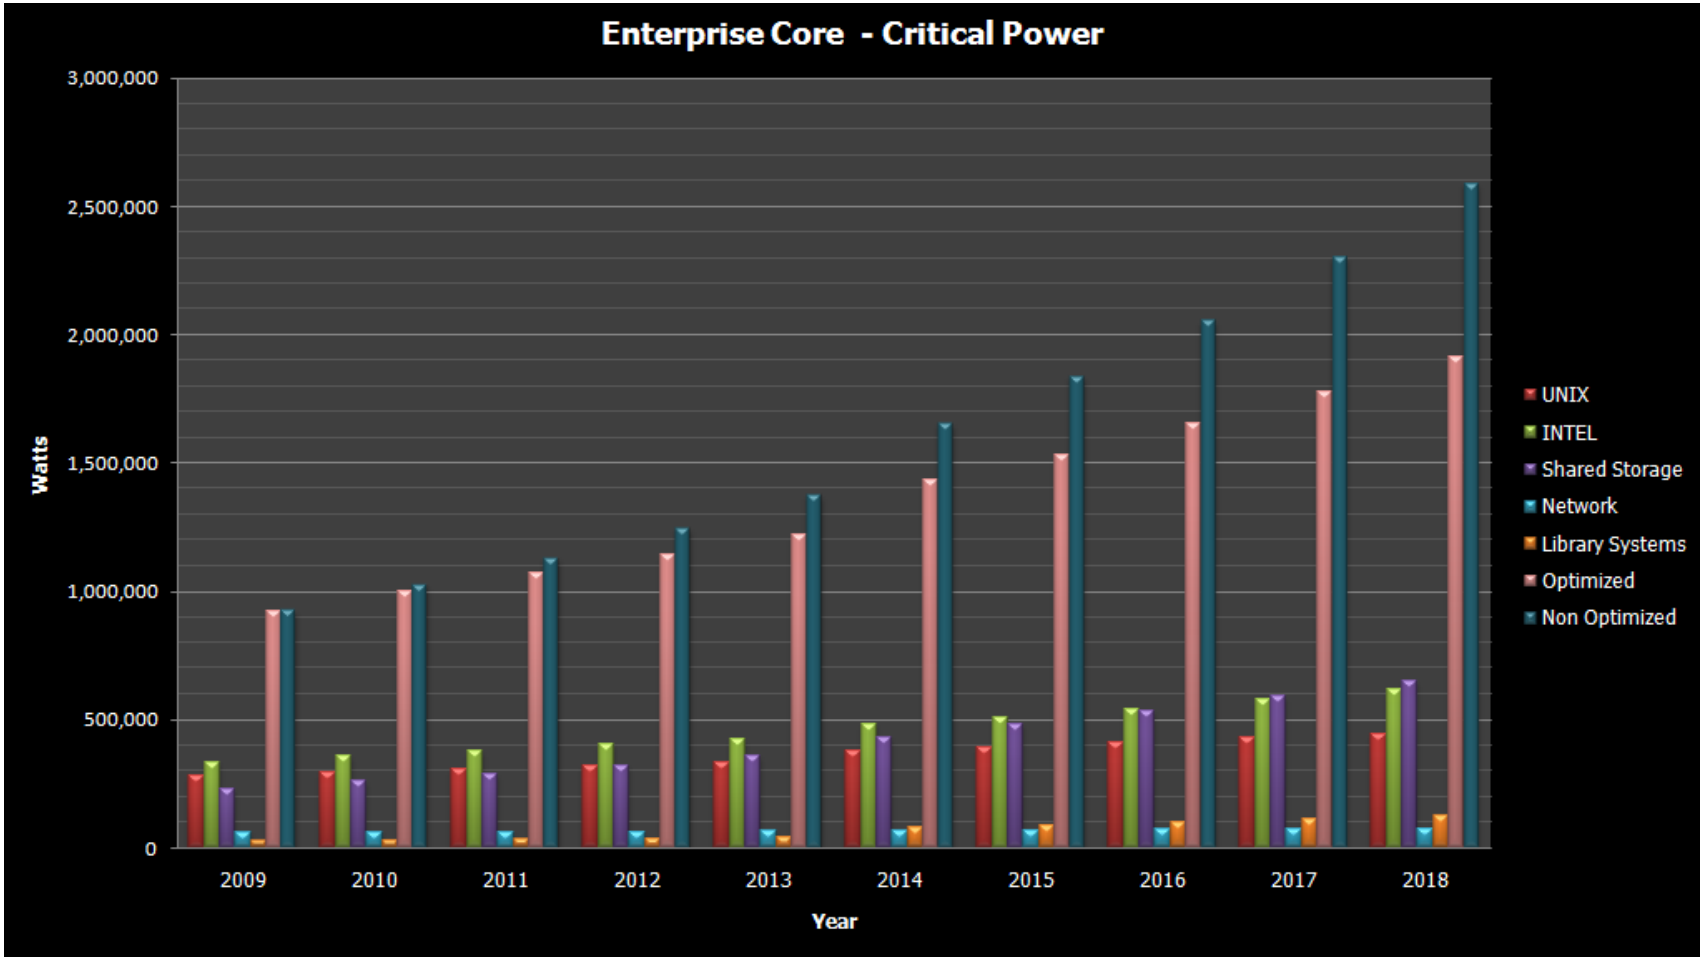

Figure 4: Ten Year Enterprise Core Critical Power Projections

- The figure above shows that the critical power requirement is approximately 1.25 mW of total power for the Enterprise Core 5 year (Day 1) growth planning.
- The figure above shows that the critical power requirement is approximately 2.0 mW of total power for the Enterprise Core 10 year (Final) growth planning.
- Any initiatives that could increase or decrease the technology platform will change the metrics accordingly. For instance, if University of Utah deploys server growth in lieu of blades, the power requirements will change over time.

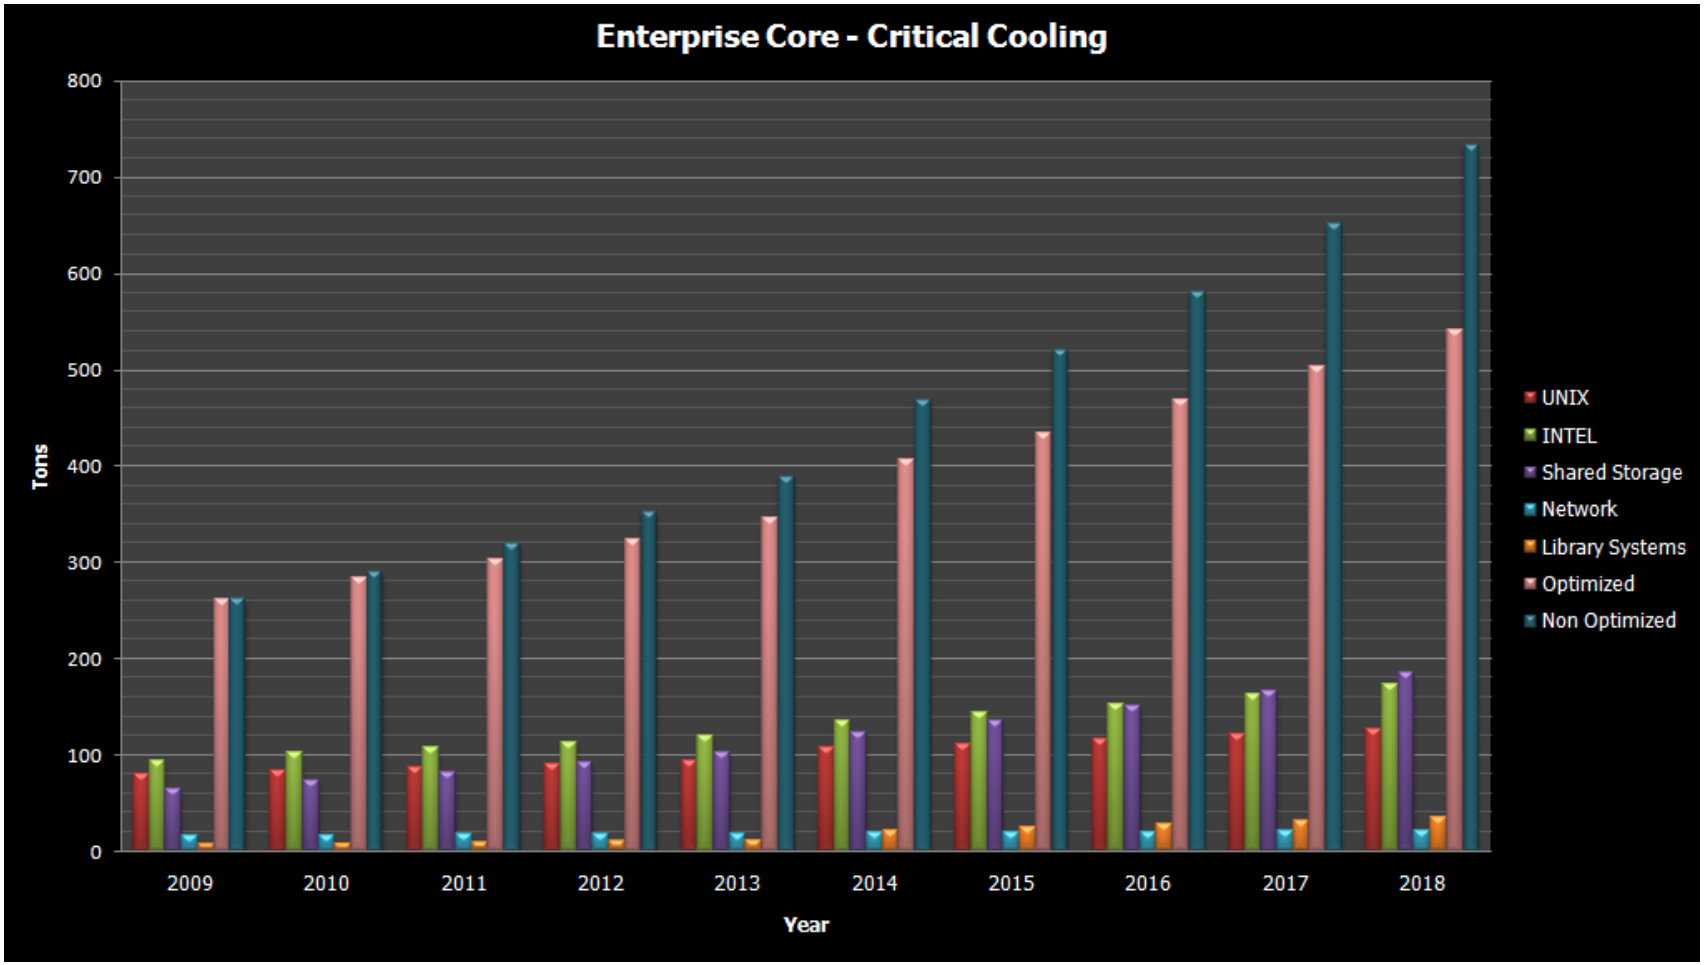

Figure 5: Ten Year Enterprise Core Critical Cooling Projections

- The figure above shows that the critical power requirement is approximately 350 Tons of total power for the Enterprise Core 5 year (Day 1) growth planning.
- The figure above shows that the critical power requirement is approximately 550 Tons of total power for the Enterprise Core 10 year (Final) growth planning.
- Any initiatives that could increase or decrease the technology platform will change the metrics accordingly. For instance, if University of Utah deploys server growth in lieu of blades, the power requirements will change over time.

3 Center for High Performance Computing Concepts

3.1 Overview

Strong support to develop a dedicated space for high performance computing was communicated throughout the stakeholder meetings. EYP MCF recommends building a dedicated space for high performance computing and believes it is in University of Utah's best interest to do so. EYP MCF believes separating the HPC area from normal enterprise type computing is a University best practice.

EYP MCF has interviewed many research primary investigators and other research thought leaders. Inputs from other universities are contained in the following summary statements:

- New grant applications that requires computational systems must include details of how these systems will be supported within the data center infrastructure. This is a mandatory requirement and equally weighted with the science portion of the grant application. Many research grant peer review processes have a mandatory review of infrastructure available to support grant specific systems.
- Grants are subject to audits that determine if the implementations of computational systems meet grant application commitments.
- General concern from research communities that major universities communities will not be able to build computing infrastructure fast enough to meet future research computational requirements.

The statements noted above apply to University of Utah as well. Additional major drivers supporting this decision to develop dedicated HPC specific space within the data center are:

- There will be a significant challenge to meet even short term growth (1-3 years) of HPC systems within the current data center infrastructure.
- Modern HPC computing requires high capacity power and cooling systems not commonly in use in business and academic data centers.
- The overall future state growth projections for HPC systems indicate that current data center facilities will not be able to support long-term growth of these very important systems.
- Research institutions are all struggling to meet the high growth and power/cooling density demands of modern HPC systems. EYP MCF believes that purpose built HPC data centers or separate spaces within the data center are the best option to meet these growth demands.
- No known solution/plan exists to accommodate long term growth of HPC/Research specific systems.
- Concern that research computing has already outgrown current data center support capacities.
- Consensus that average life span of research systems is approximately 3 years.

Quality of research computing facilities is increasingly becoming a point of separation for top institutions. Existing data centers cannot accommodate the projected high growth rates, especially for power and cooling. New researchers often require research computing resources immediately and top researchers also bring funding opportunities if computing facilities are available. This highly dynamic growth pattern is VERY difficult to predict and purpose built research data center environments cannot be quickly stood up due to very high power density requirements. This has led many leading research universities to adopt a build it and the grants will come approach.

EYP MCF believes that all major research universities that wish to remain at the forefront of their peers must move to an on-demand type model for their research communities. Leading research universities must provide a cost effective ability to provide on demand data center space, power and cooling capabilities for rapidly changing research computing demands.

EYP MCF believes that the universities that find a way to invest in their research future will continue to be the premier universities of the future. This capability will attract and retain the type of research staff that leading research universities demand to maintain their leadership in excellence.

The consensus recommendation between University of Utah Stakeholders and EYP MCF is to design and build a separate isolated HPC/Research specific area within the data center.

3.2 Physical Data Center Planning – HPC

Based on input from CHPC stakeholders, the Center for High Performance Computing Area should contain the following attributes. EYP MCF discussed in detail the needs of CHPC and supports their estimate for space, power and cooling needs are not unreasonable and fall in line with other university programs.

HPC Area

- 4,000 sq. ft. raised floor (Day 1)
  - potential to expand to 8,000 sq. ft. (Final)
- 4 foot raised floor preferred
- Hot aisle containment capability
- Non-redundant CHW distribution system below raised floor – for future use
- 1mW critical load (Day 1)
  - ability to expand to 3mW (Final)

HPC Core – enterprise type systems for CHPS which require similar infrastructure redundancy as "Enterprise Core"

- 1,000 sq. ft. raised floor (Day 1 only – does not expand)
- Tier 2 infrastructure
  - N+1 UPS
- Utilizes Enterprise Core Infrastructure – 75kW Day 1 restriction

4 Co-Location

4.1 Overview

Co-Location offering is a concept the University is interested in offering colleges and other entities with University relationships. At the time of this study, specific needs for a dedicated Co-Location offering had not been quantified. Due to the lack of quantifiable interest in external parties, and initial cost estimates of allocating space specifically for Co-Location offering, the decision was made by University of Utah stakeholders not to include a Co-Location offering in the initial build data center program.

The discussions contained in this section are only initial thoughts and discussion points for a Co-Location offering. Once the decision was made to not include the Co-Location within the initial build program, no further evaluation was conducted.

The possibility of new University operated Fiber Optic MAN opens the possibility for a Co-Lo offering. The initial concept is to offer colleges and other entities with University relationships a Co-Location offering in a new hardened data center. The space would "cage off" and separated from the Enterprise Core. A "meet me" room would be used to hand off carrier connections.

The following is a listing of entities considered for consultation to determine the interest level for a Co-Location offering.

Potential University Departments	Potential Colleges	Potential entities outside of U of U
• Ucard, bookstore • GIS • Development Office • Project Management Group • Facilities Management Group • Campus Design Construction • Planning • Operations • Student Systems • Financial Imagining - FORTIS	• Law • Fine Arts • Architecture • Business • Behavioral Sciences • School of Medicine • College of Nursing • College of Pharmacy • College of Health • Mines and Engineering • Education • Social Work	• ARUP labs • Intermountain Healthcare • State board of regents • Utah State - HPC • Southern Utah - HPC / DR • Dept of Technical Services • State Dept of Health • Medical Collaboration • American Geological Institute • HHMI • Brain institute • USTAR • EGI Energy and Geophysics • Scientific Computing Institute • Huntsman Cancer • Conflict of Interest

4.2 Physical Data Center Planning – Co-Location

Based on input from University of Utah stakeholders, the Co-Location Area should contain the following attributes. EYP MCF discussed the anticipated needs of a Co-Location offering and believes their estimate for space, power and cooling needs are not unreasonable if the prospective interest is accurate.

Co-Location Area

- 5,000 sq. ft. raised floor (Day 1)
  - Potential to expand to 7,500 sq. ft. (Final)
- 150 w/sq. ft. – 750 kW critical load (Day 1)
  - Ability to expand to 1 mW (Final)

To date none of the outside entities polled have expressed any desire to collocate their IT systems in this new data center.

5 Data Center Program Recommendation

This is a joint recommendation of both EYP MCF and the University of Utah IT steering committee.

5.1 Overview

The University has already identified the need to implement a new purpose built data center and has already taken the initial step of procuring an existing structure located within the Salt Lake City area at 875 S. West Temple Street. EYP MCF supports the decision by the University and agrees the current data center topology will not meet the University's future needs and a purpose built data center is needed. The program recommendation outlined below supports the Day 1 needs as reflected in modeling for consolidation of the seven identified data centers into a single data center.

With the stated requirement that 24/7 access to systems is becoming the required norm by University of Utah customers, and hosting of hospital systems, a concurrently maintainable data center is recommended (Tier III). Bandwidth capabilities need to be investigated to ensure DR and HA scenarios can be supported by current wide area network infrastructure.

The Day 1 needs of both Enterprise Core and High Performance Computing can be met by locating the data center within the "south bay" of this existing structure.

Enterprise Core Program Build Concept

- Accommodates Enterprise Core Program Build needs and removes the corresponding loads from existing datacenter spaces. (refer to Section 2 - "Enterprise Core Overview")
- Richfield will continue to be disaster recovery site.
    - Richfield may provide the ability to house some High Availability (HA) systems. HA systems would be split between two data centers.
- Existing ACS and Marriott Library data centers may be utilized for HA scenarios.
- Dedicated operations staff with limited access to data center floor.

HPC Program Build Concept

- HPC unique requirements are met with infrastructure designed specifically for HPC computing.
- Accommodates CHPC Program Build needs and removes the corresponding loads from existing datacenter spaces.
- HPC Core area will utilize Enterprise Core infrastructure.
- Separate entrance and physical isolated from Enterprise Core to accommodate less stringent access requirements.

Co-Location services are **NOT** included in the initial Program Build data center plan.

5.2 Opportunities and Challenges - Program Recommendation

Opportunities	Challenges
Build a new data center with up to date redundant infrastructure to support concurrent maintainability and utility power interruptions.	EYP MCF believes the University of Utah is at risk of increased frequency of data center outages due to current facility infrastructure capacity/reliability issues.
Build a new consolidated data center with the capacity to accommodate anticipated growth projections. Implement an operations model which allows limited controlled access. Continue to implement optimization methods such as virtualization technologies.	Capacity issues will limit the ability of the University to support continued development of operational systems.
Build a new consolidated data center with the capacity to accommodate anticipated growth projections. Implement an operations model which allows limited controlled access. Continue to implement optimization methods such as virtualization technologies.	Current data center(s) capabilities cannot sustain University systems growth projections.
Continue to encourage vendors to support virtualization technologies. Evaluate possible competitor's virtualization compatibilities and support.	Lack of vendor support for virtualization in both Enterprise and Healthcare systems is limiting consolidation efforts.
Build a space within the new data center dedicated to high performance computing unique infrastructure requirements.	The risk of not being able to compete for grants due to inability to meet computing requirements.
Evaluate current bandwidth capabilities of existing Wide Area Network. Continue to pursue current initiative of University operated MAN which would provide access to Richfield.	Would need substantial bandwidth from campus and Richfield to new data center to accomplish intended service offerings, disaster recovery and high availability.
Build a space within the new data center dedicated to high performance computing unique infrastructure requirements.	High performance computing systems have unique requirements compared to University core systems and is best served in a separate environment designed to different reliability standards.
Significantly reduce the cost to operate and maintain by utilizing a single large data center within the University of Utah environment.	Consensus must be gained by current data center customers that remote access to systems is acceptable.

5.3 Physical Data Center Planning – Program Recommendation

Program Recommendation Concept - accommodates Enterprise Core and CHPC Day 1 needs

- Enterprise Core Area
  - 11,000 sq. ft. raised floor
  - 1.25 mW
- HPC Area
  - 4,000 sq. ft. raised floor
  - 1 mW
- HPC Core - requires similar infrastructure redundancy as “Enterprise Core”
  - 1,000 sq. ft. raised floor
  - Utilizes Enterprise Core Infrastructure – 75kW Day 1 restriction

5.4 Program Recommendation – Test Fit

The objective of test fit below is to demonstrate a general configuration of white floor space layouts. The overall intent is to represent the types and quantities of equipment that could be placed in the white floor space, and is not meant to be a recommendation for final configuration.

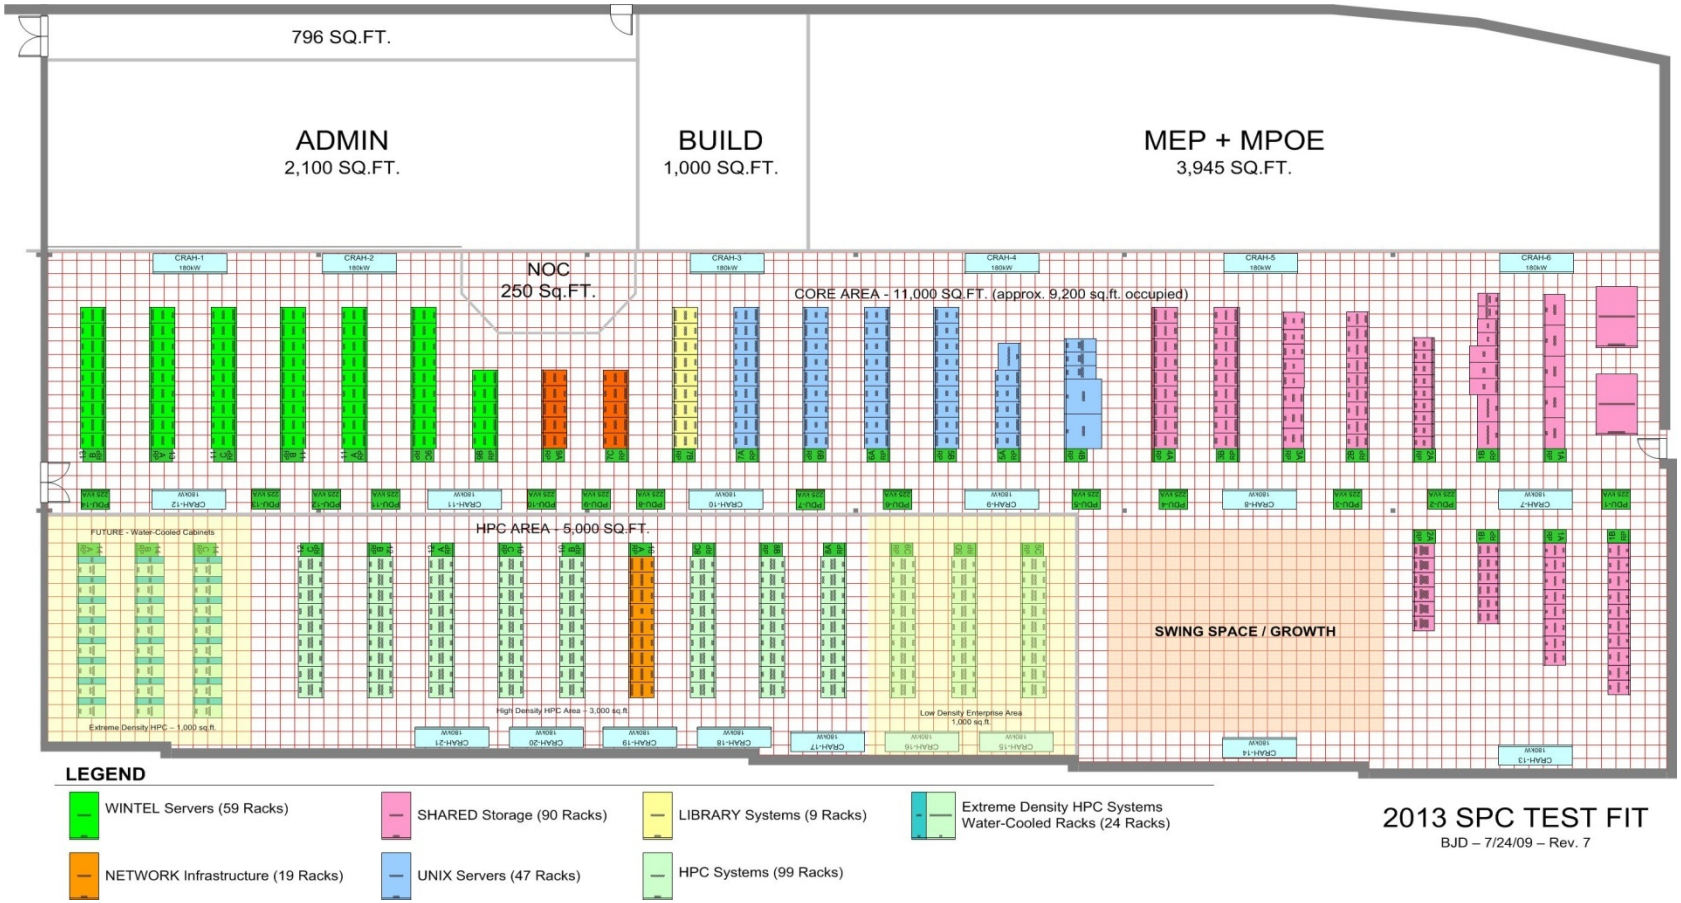

- Core Area: contains the anticipated cabinets for Day 1 (2013) based on space, power and cooling modeling conducted for the Enterprise Core Area.
- HPC Area: contains a representation of quantity of racks which could possibly be located in this space. Specific modeling was not performed for the HPC Area.

Figure 6: Program Recommendation – Test Fit

5.5 Mechanical and Electrical Concepts

See the separate document titled “University of Utah Solicitation 9967 - Data Center Improvements – Programming Basis of Design” for Mechanical and Electrical Facility Concepts.

5.6 Program Recommendation ROM Cost Estimate

All construction cost estimates were developed using the Uptime Institutes (TUI) cost estimator that calculates the MEP infrastructure cost based on critical KW of power utilized. Added to the cost of the MEP infrastructure is a base cost of $150 / square foot for the data center building shell. This cost was reduced to half of the $300/sq.ft. amount recommended by TUI to account for shell of building existing. EYP MCF assumes no responsibility for the accuracy of these cost estimates. They are solely based on the Uptime Institutes cost estimator model.

Costing assumptions utilize TIER 2 data center design concepts for Enterprise Core and Tier 1 sizing concepts for HPC as defined by the Uptime Institute. These are ROM costs for comparative purposes only and not intended to take precedence over any cost estimates provided by Skanska for this project. Final Option is a ROM estimate as if doing full build today.

Summary Comparison of Options
Rough Order of Magnitude Cost Estimate (±30%)
Does NOT include COLO

	Final Build EC & HPC Tier 3 Option	Final Build EC & HPC Tier 2 Option	Program Build Tier 3 Option	Program Build Tier 2 Option
White Floor Space	25,000 sq. ft.	25,000 sq. ft.	**16,000 sq. ft.**	**16,000 sq. ft.**
Critical IT Load	6 mW	6 mW	**2.25 mW**	**2.25 mW**
Cost	$84,100,000	$63,100,000	**$42,650,000**	**$29,525,000**

These cost estimates were utilized during the program development for budgetary guidance only. Please refer to the construction cost estimate for the specific cost estimate developed for the recommended program build details.

6 Best Practices

6.1 Critical Issues Best Practices Gap Analysis

The following outlines core technology critical issues discussed during the stakeholder meetings with any deficiencies (gaps) discovered during the stakeholder interview process.

Disaster Recovery / Business Continuity Planning

University of Utah Current State:
University of Utah does not have a well defined and executable disaster recovery / business continuity plan. Specific university wide plans were not shared with EYP MCF. All current recovery planning and High Availability capabilities for computing environments are developed and implemented at the local department and/or college/campus level.

Best Practice:
A DR/BCP plan that clearly identifies the Recovery Time Objectives (RTO) and Recovery Point Objectives (RPO) for both data and applications.

This plan should be exercised at least once each year for all critical University of Utah application environments. This exercise should be monitored by the business units impacted and with lessons learned applied to gaps discovered during the exercise.

EYP MCF Gap Analysis:
EYP MCF has found that University of Utah does not meet the normal level of preparedness for entities of University of Utah size/revenue.

It is the EYP MCF opinion that the current level of DR / BCP planning would prevent University of Utah from successfully recovering their business in the event a catastrophic site event occurred at University Park. An event impacting the entire University Park campus would obviously need to be a large scale event, but if a Pandemic Event occurred other University of Utah campus sites would not be able to function on many critical levels.

Application Criticality Prioritizing Organized by Business Unit Requirements

University of Utah Current State:
EYP MCF did not find consistent application priority definitions within the University of Utah environment. Various local departments and/or college/campus levels have taken measures to back up critical data at offsite locations. However, this is not the case for all levels throughout University of Utah.

Best Practice:
All applications within the University of Utah IT environment should be classified by criticality within a Priority Application Range. Typically applications are organized by Priority designation of Priority 1 through Priority 4 with Priority 1 applications being the most critical to the business. RTO's and RPO's should be defined and communicated by both IT and business unit managers for each defined application Priority type.

EYP MCF Gap Analysis:
Lack of any specific application prioritization increases the probability that University of Utah cannot successfully recover from a catastrophic site disaster event regardless of time passage from the event.

It is EYP MCF's opinion that if University of Utah does not understand the criticality of applications based on business requirements then the chances that any successful recovery of the production IT environment would be problematic at best and impossible at worst if a Primary Data Center site wide disaster occurs.

Data Replication and Storage Standards

University of Utah Current State:
Minimal data replication and storage standards have been implemented within the University of Utah IT infrastructure.

Best Practice:
Application data stores associated with Tier 1 and Tier 2 application environments should be stored on segmented SAN type storage devices. Lower Tier application environments data may be stored on lower cost separate SAN/JBOD type equipment.

Data storage technologies should be ranked in a Tier system that defines levels of redundancy and reliability corresponding to the application Tiering systems. As an example Tier 1 data storage would map to Tier 1 application data stores and provide the highest levels of redundancy and reliability utilizing technologies such as RAID (5,10, etc.) and hot spares. All Tiers of applications should have a corresponding data storage Tiers specification that defines the redundancy/reliability required to meet the specific RPO/RTO for each application Tiering specifications.

The Tier 1 or Tier 2 disk based data storage is replicated either asynchronously/synchronously to a site that is at most 40 kilometers remote to the primary production data center site. EYP MCF does not recommend attempting to synchronously replicate data a distance greater than 40 kilometers.

Deployment of D-Duplication Technologies to minimize duplication of data is recommended.

EYP MCF Gap Analysis:
University of Utah should start a process to indentify data replication and storage standards for at least Tier 1 application environments.

High Growth Area – High Performance Computing (HPC)

University of Utah Current State:
University of Utah stakeholders identified that HPC has been a major contributor to data center capacity issues. EYP MCF has found that HPC is a high growth area for almost all universities

Best Practice:
Provide a data center space designed specifically for high performance computing.

This allows for implementing advanced infrastructure designs unique to HPC while taking advantage of the typically lower redundancy requirements for facilities.

EYP MCF Gap Analysis:
EYP MCF spoke with University of Utah stakeholders at length to understand the effect HPC has had on the current University of Utah data centers. The stakeholders interviewed expressed concern that research computing has already outgrown current data center support capacities.

There was consensus that average life span of research systems is approximately 3 years. EYP MCF therefore recommends building a standalone purpose built HPC specific data center space designed to meet the very specific demands of HPC

type computing. It is very difficult to collocate HPC type computing systems with business and academic systems within the same DC space/environment.

The overall future state growth projections indicate that current data center facilities will not be able to support long term growth of these very important systems. Removing these systems from Enterprise Core space will help extend the life of Enterprise Core space.

IT Systems Deployment Standards

University of Utah Current State:
University of Utah has defined IT architecture deployment standards that are based on current production data center infrastructure capabilities. EYP MCF has found that these deployment standards are followed and well implemented in the current data centers. An area that may be lacking, but being addressed at individual organizational levels, is in the High Availability environment.

Best Practice:
Well defined IT architecture deployment standards for the following areas:

- High availability
- Virtual servers
- High density systems (blade chassis)
- Medium density systems (rack mounted servers)

EYP MCF Gap Analysis:
University of Utah is utilizing the existing data centers to the fullest capacities possible with their existing constraints. However, standardizing throughout the University on platforms and virtualization, better utilization of university resources as a whole can be achieved through resource sharing.

6.2 IT Operations Best Practices Recommendations

EYP MCF has found that several IT operational enhancement efforts are defined and in process within the University of Utah global IT organization. EYP MCF has reviewed the ongoing operational efforts and provides the following recommendations to supplement these efforts:

- Continue with all initiatives to virtualize physical servers and develop standardization on one virtualization platform. The ratio of virtual instances to physical servers of 10:1 should be a minimum implementation ratio. Immediate focus should continue to be placed on virtualizing new servers that are placed into production. EYP MCF recommends that University of Utah adopt an application deployment policy that states all NEW application servers must be implemented via virtual server technologies unless it can be proven that the new application environment will not work adequately in a virtual environment. An exception to the virtualized server standard deployment rule would have to be granted by University of Utah executive management.

- Develop a standard technology refresh program that is designed to refresh older technology systems on a reoccurring time line. EYP MCF has found that an annual technology refresh review is required to provide adequate planning for this important operational requirement.
    - Develop optimized cabinet layout/deployment standards through planned future state growth efforts. The current recommendation to implement a new Data Center within a new facility is an ideal opportunity to update and enhance current cabinet deployment standards.
    - Define critical applications priority matrix cross referenced by associated computing hardware. EYP MCF has found that the University of Utah IT organization has not developed a critical application Priority plan that define ALL applications based on business unit level criticality requirements. This information is crucial to the ability for any business to survive any significant catastrophic event. In addition to Prioritization, University of Utah should define the Recovery Point Objective (RPO) and Recovery Time Objective (RTO) with cost of downtime for each application.

6.3 Limit Physical Access to the Production Data Center

A critical data center best practice is removing people from the production data center whenever and wherever possible. It is a well-proven fact that direct human error is responsible for almost 50% of all unplanned outages within production data centers.

Modern data centers are expensive purpose built facilities expected to provide high availability. Locating offices or providing working areas for people within the data center is not a cost effective use of this expensive space. For maximum protection against unplanned outages, limit personnel access into the data center proper and all support equipment spaces. Require escorts for all support equipment service personnel.

Provision the data center to allow remote management and administration of IT systems. Most IT systems delivered within the past three years feature remote management capabilities via a network connection. If remote management capabilities do not exist, cost-effective network based access solutions are readily available on the commercial market.

All current clients of EYP MCF provision new data center spaces for minimal people areas and adopt a lights out operational model. In our experience, other Universities planning to build new data centers are adopting this minimal people footprint operational model. EYP MCF recommends that University of Utah adopt this best practice.

6.4 Consolidation of Multiple Smaller Data Centers into Larger Data Center

The electrical/mechanical infrastructure costs constitute approximately 75% of the overall costs of data centers. Building higher numbers of smaller data centers leads to the trend to place excess electrical/mechanical capacities at each data center for growth requirements.

A simple cost analysis can be provided by looking at how electrical capacity requirement drive data center costs. Using current day Uptime Institute published cost estimation data, we can assign a cost of $12,500 per KW of UPS systems for a TIER 2 facility. Additionally, $300 per Square Foot (SF) of raised floor shell cost must be added to the per KW cost.

Smaller Distributed Ad-hoc Data Center Cost Scenario

Assume 8 data centers with following attributes:

- TIER 2 Implementation of Mechanical/Electrical/Plumbing (MEP) Infrastructure
- 2000 SF of raised floor area.
- 250 KW of redundant UPS capacity (175 KW/site utilized – 1.4 Megawatts total all 8 sites).
- 75 KW of standalone UPS growth capacity at each site – cannot be shared!
- 125 Watts per SF MEP capacity.
- 70% utilization of UPS capacity on average between all the sites.

Associated costs for each data center:

- Overall cost to build each facility – ($12,500 X 250) + (2000 X $300) = $3,725,000
- Excess electrical capacity of each data center (average) = 75KW
- Excess total electrical capacity of all DC's = 600KW
- Cost associated with 600KW of unutilized TIER 2 electrical capacity = $7,500,000

- MEP Build costs of all 8 data centers - $29,800,000.
- Must actually build more capacity at any single site that requires more than 75KW of additional electrical capacity.

By utilizing several smaller data centers the excess electrical capacity is NOT accessible to all groups within the University of Utah system.

Larger Single Data Center Cost Scenario

Consolidated single data center with conservative space, power and cooling optimization applied:

- TIER 2 MEP Infrastructure Implementation.
- 1 data center with 14,000 SF of raised floor area (utilizing conservative consolidation optimization).
- 1.75 KW of redundant UPS capacity (1.4 Megawatts Used).
- 350,000 Watts of UPS growth capacity - available for ALL University of Utah user groups.
- 125 Watts per SF electrical / mechanical capacity.
- MEP Build cost for data center – ($12,500 x 1750) + (14,000 x $300) = $26,075,000
- 80 % utilization of shared MEP infrastructure. Better use of expensive MEP infrastructure.

Cost Comparison

The MEP infrastructure build out cost savings for this simple comparison is $3,725,000 alone. This savings equates to 12.5% reduction of building 8 data center vs. 1 data center with very conservative consolidation optimization.

Overall UPS electrical capacity IS the primary cost driver for data center spaces. It is extremely important to optimize the utilization and flexibility of expensive MEP infrastructure capabilities.

The operational cost savings by consolidating smaller data centers into fewer and larger data centers can be conservatively estimated at a 10% savings as well. EYP MCF has seen much higher cost savings from real world consolidation implementations. Nevertheless, for the purpose of this project EYP MCF feels that a 10% cost IT and facilities cost reduction goal is a very conservative value to utilize.

Efficiency of Utilization

The single site consolidated data center provides maximum flexibility to meet new loads from any group utilizing the data center. Also adding MEP infrastructure capacity at the consolidated data center provides this expanded capacity for all University of Utah user groups, not just a single small data center site.

6.5 Recommended Next Steps

EYP MCF encourages University of Utah to engage in a series of initiatives that will reduce risk and provide a stable environment for migration into a new Primary Data Center and to support future growth. Implement these initiatives using a multidimensional approach and execute in concert with each other. The order does not imply any priority or precedence of one over another.

The initiatives are as follows:

- Continue to pursue optimization efforts as become available.
- Consideration of best use for Park Building and Marriot Library for future state High Availability.
- Implement a project to determine what applications require deployment in a High Availability application environment. Limit HA application deployments to only the applications that are deemed business critical to University of Utah.
- Continue with Storage Tiering, DDUP and Thin Provisioning for disk based data to reduce storage growth.
- Evaluate current bandwidth capacity to Richfield and possibility of University owned MAN.
- Implement a project to define operations methodology at new data center.
- Begin marketing the advantages of a single centralized data center to the owners/operators of the many independent server rooms distributed around campus.
- Begin development of a migration plan. Moving hardware without disruption of computing resources will require detailing planning. There are several questions besides the obvious concerns of what order to move equipment and when can the respective users tolerate the disruption of service.
    - What equipment can the University eliminate prior to move-in?
    - What equipment can the University consolidate into more energy efficient hardware?
    - Which applications have porting constraints that would preclude migrating to newer, more efficient hardware?
- Begin investigating cabinet vendors to select a standard rack mount style cabinets and cabinet-PDUs for the new data centers. See EYP MCF best practice recommendations for cabinets.

7 Data Center Master Planning Methodology Discussion

7.1 Future State Space, Power and Cooling Requirements

Ten Year Space, Power and Cooling Requirement Estimates and Assumptions

EYP MCF developed a ten year consolidated IT systems Space, Power and Cooling (SPC) requirement projection for five conceptual scenarios based on input received from University of Utah stakeholders. The main input used to model the ten year conceptual scenarios SPC estimates were equipment inventories collected during EYP MCF C&I infrastructure discovery and/or estimates provided during stakeholder interviews. It is important to note that the accuracy of these estimates significantly decrease after five years, as it is nearly impossible to predict technology advancements and their effects on the data center past five years.

The overall factors accounted for in the SPC calculations are as follows:

Windows

- OIT – Growth rate of 30% last 12 months (recent Messaging upgrade).
  - average of 30% growth every 3 years
- ITS – 10:1 Virtualization Ratio on blade systems
- ITS – virtualization is leveling off
  - growth of 2-4 enclosures per year = 1 rack every years
- ACS – current Virtualization Ratio is 15:1; Future to be 20:1

UNIX

- OIT/ITS – 25-30 T2000's Day One
  - 50:1 virtualization
- UEN – Virtualize SUN E6000 to 16:1 ratio
- ITS – records indicate 10% growth
- ACS – 10:1 virtualization rate

Storage

- ACS - Current Hitachi Storage will reach capacity within 12 months
  - = 3 year / system growth rate
- UEN and ITS – data doubling every year
- OIT – Significant Growth to support media streaming, virtual machine images, museum virtual tours
  - NetApp 3020 to support growth

Applications

- Large growth of Citrix Servers
- Data Warehouse driving doubling of storage yearly
- PACS imaging causing extreme growth in data storage requirements
- New daybreak clinic may add spike in growth

Space Planning Terms and Definitions

Two major planning factors are included within the EYP MCF space planning modeling process. These two major planning factors are:

a) **Grossing Factor** – The grossing factor defines all of the space not specifically associated directly with IT equipment within the raised floor area. The space planning process used the specific IT equipment list provided by University of Utah stakeholders to develop a current state baseline. This IT equipment only baseline models the space required using current state best practice space planning methodologies. The model applies the grossing to the IT equipment only space requirement to determine the total raised floor area required. The grossing factor includes space for at least four-foot hot/cold aisles, PDU's/RDP/CRAC units on the raised floor and other important non-IT equipment specific spatial requirements for a normal raised floor area. The grossing factor provides for efficient raised floor area modeling without being overly aggressive or conservative for future state spatial modeling.

b) **Swing Space Factor** – The swing space factor adds a minimal amount of raised floor space onto the final year projections of total raised floor area. The swing space factor only adds relatively low cost raised floor area (does not include adding the higher cost MEP infrastructure) only as a long-term space planning safety factor. The concept of swing space would allow staging new technology for testing at the end of the ten-year planning cycle. After placing the new technology into production, remove the existing technology from production to eliminate the associated MEP loads.

7.2 Data Center Planning Considerations

Overview

The section that follows briefly describes some of the more critical aspects for the University of Utah's ITS department to consider during development of best practices for their data center and the University of Utah's Enterprise in general. Every Enterprise has different requirements to support their business, so these recommendations are general and are not University of Utah specific.

The IT organization, as well as the Executive level, must provide a solid commitment to ensure the success of those individuals tasked with deploying standardized best practices. This includes necessary funding, time, and labor to complete the tasks.

Standard Operational Procedure (SOP) is critical for every organization. Essentially a SOP describes the organization's approach for performing a major task (i.e. deploying a server). Within every SOP are several sub-tasks, or methods of procedure (MOP), required to complete the major task. A MOP includes the detailed instructions for each minor sub-task required to support the SOP (i.e. load server OS, configure server OS, obtain network address, obtain DNS name etc…). Some organizations do not require a detailed SOP, or MOP, for every item and task, but have defined standardized procedures and methods, which are critical for efficient deployment and support.

Many organizations take their best practices quite seriously because they understand the benefits of the best practices model, as well as the detrimental side effects if they do not define and follow operational procedures. Some organizations have a military approach to control of these procedures, with consequences tied to poor adherence of the Executive supported efforts. Executive support helps enforce acceptance and compliance of the organization's best practice policies.

Standardization

Standardization should be one of the first goals for any IT organization. This can apply to procurement, support, warranties, platforms and procedures. This is critical in Engineering of future solutions as well as the deployment and support of the platforms that the solutions depend upon.

The bulk of services deployed will satisfy the requirements set forth by the organization. It is also important to understand that there are the occasional systems that require a modification to the deployment and support model. The organization should assess these "one-offs" on a case-by-case basis and try to limit them as much as possible.

Deployment Model

Standardizing the deployment model for an organization produces a baseline and defines how systems are deployed and managed. This model provides the underlying foundation for the whole facility and the services within. Typically an effort of this

size would start at the site level, where the base infrastructure resides, and be taken through the various components of the data center such as power, cooling, physical infrastructure, cable plant, network, servers, storage, and applications.

Once the various primary support components are clearly defined, as well as the individuals responsible for those components, then an outline can be written defining areas to be addressed. Through a methodical and detailed analysis, documentation can be developed based on existing operational procedures. Typically the initial approach in defining a standardized deployment model is broken out into three or four primary technical areas, for example:

1) Physical Components
2) Logical Components
3) Support and Operations

Once the primary areas are defined, then the multiple sub-areas should be outlined, and the procedures for each sub-area documented. Due to the nature, size, scope and complexity of IT systems and data centers in general, these recommendations are provided as an outline for University of Utah to consider during the development of their internal best practices and procedures.

Once all procedures are documented, typically a technical review board is established to proof and edit all of the proposed procedures for the deployment model. This board can also be responsible for ensuring that a holistic approach is taken, and that the proper overlap is provided for thus ensuring that no technology or process gaps exist.

The review board should think in terms of a layered approach and have representation from various technical disciplines. In the future, members of the review board can take part in weekly change management meetings to ensure that operational continuity, best practices, that the deployment model and procedures are followed. Based on the three aforementioned technical areas, a number of best practices and considerations are provided below.

1) <u>Physical Components</u>

a) ***Power*** – All circuits, RDP's, PDU's, breakers, power strips and power cords are properly labeled and monitored. All whips and power cords are of the appropriate length without excessive slack. All power cords in the cabinet provide enough length to support the device and that excessive slack does not overcrowd the rear of the cabinet when dressed. All power elements have SMTP notification; IP based alarming and monitoring capabilities.

b) ***Cabinet Space*** – All cabinets have a standardized configuration based on cabinet type (network, server, etc…). All cabinets are labeled front and rear and have legible RU elevations marked on the rails. Adequate space is provided in the front and rear for vertical cable management. Cabinets have proper airflow and are positioned in a hot/cold aisle

configuration. Any feeds from below the raised floor should have openings only as large as required and have some type of device that minimizes air bypass through each raised floor opening. All cabinets are properly grounded.

c) ***Cabling*** – A structured cable plant provides great flexibility for service and operations as well as the proper physical infrastructure to support the deployment model. This design revolves around the needs and requirements defined by the previously defined network, SAN and server deployment models. The structured cable plant should be scalable and allow for current and future transmission technologies. Various other items should be included in the design such as pathway redundancy, infrastructure support/pathways and bonding/grounding of the various infrastructure elements. Adequate fiber and copper should be in place to support network and SAN requirements. The cable plant infrastructure should provide room for growth and support the transmission of 10 Gigabits per second. All patch panels, optical panels, ports and cables should be properly labeled and documented. All patch cords should be properly sized and of the proper length, without excessive slack, dressed neatly and labeled. The entire cable plant should be fully documented with link and channel information. The cable plant documentation should be updated on a consistent basis and have a pre-defined labeling scheme. This labeling scheme should be pre-defined and is adhered to during the lifetime of the facility. This labeling scheme provides ease in troubleshooting, deployments and operations. All labels should be produced via a labeling machine (not hand written).

d) ***Data Center Access*** – Access to the facility is limited only to the people required to support the facility. All access in and out is documented and proper access controls are in place and logged for future reference. Cameras or other monitoring devices provide historical access data for the facility.

e) ***Environmental*** – All cabinets (or data center zones) are monitored for temperature and humidity levels. The proper fire protection is in place and has been tested. All environmental elements have SMTP notification, IP based alarming and monitoring capabilities.

2) <u>Logical Components</u>

a) ***Network*** – University of Utah should standardize on Network hardware platforms and IOS versions to ensure ease of management. It is wise to maintain cold spares on site for the most critical network elements in case of hardware failure. Warranties for all components, along with the appropriate response time, to meet the business need are recommended. Network devices should have standardized configurations and are automatically backed up on a regular basis. The appropriate network based Security systems should be in place (firewalls, intrusion protection, etc…). Redundant power supplies should be used and load balanced. All power cables should be labeled indicating which power supply, power strip (A or B), whip and PDU/breaker ID.

b) ***Server Platforms*** – It is important to limit the amount of server platforms within the data center for a multitude of reasons. One of the primary reasons is to ensure that the appropriate space, power and cooling that has been provided

for in the model can be supported in the data center. It is important that facilities can accommodate the proposed deployment model on a cabinet by cabinet basis. The server model also impacts network ports, SAN ports, Management ports, cabling and the raised floor (total weight) requirements. Standardizing on server platforms also eases deployment for future applications and services. Many organizations will break this into two areas: Standard servers and blade chassis'.

c) ***Operating System*** – Standardization of Operating Systems provide for ease in deployment and management. University of Utah should have a defined methodology in place for system upgrades and patching. All versions of OS or IOS are backed up and available via the network, as well as on an alternative method (CD or flash). OS and IOS updates are checked on a regular basis, and for bugs.

d) ***Virtualization*** – University of Utah should developed a standardized virtualization model, supporting approved images (that has been properly tested) to maximize the hardware platforms supported without exceeding server or network capacity. The use of a standardized virtualization software configuration for all ESX servers and Virtual Machines is imperative to the success for this technology. Proper resource pools should be provided for to enable automated configuration. This should include that the appropriate reservations are made for DHCP/IP addresses, VLAN's, Load Balancing, Firewalls, etc...All firmware and device drivers should be consistent across all host and server platforms. SNMP should be utilized and tested to confirm that the protocol is properly and securely configured. The client should make sure that the entire VM-FS volume is backed up to the SAN on a consistent basis. Make sure that a procedure is in place for out of production network backup. Individual backup agents should be configured on the Virtual Machines and consolidated backup should occur on a physical machine or SAN drive.

e) ***Storage*** – Storage platforms should be standardized for both hardware and software configurations. Switch fabric deployments should follow network and server best practices in regards to power, cabling, labeling, configuration and labeling.

f) ***Security*** – Develop a written security policy that is approved at the CIO level, at a minimum. Compliance to this policy should be applied to every user on the organizations network. Many software and hardware based solutions are available to customize the university's security policy to a technologically based solution. The best approach to security is a multi-tier approach, which starts at the perimeter of the network and touches every element of the network. This ranges from the routers and switches down to the desktops. Hardening of server operating systems is recommended. University of Utah should utilize firewalls, intrusion detection devices and extended access-lists at a minimum. It is prudent to have quarterly, bi-annual, or annual security audit/s from an external organization that can try to penetrate various resources via both internal and external methods. This audit should be kept quiet from the user population to ensure that the penetration testing is conducted during normal operating conditions. There are too many items to list in this forum.

g) ***Application*** – Application best practices truly depend on the application being used and the environment in which it resides. Seek information from the application vendor or developer and adhere to the minimum platform recommendations within this document (server, network, security, SAN, etc…).

3) Support and Operations

a) ***Deployment Procedures*** – Standardizing the deployment procedures for all classes of devices eases the actual procedure and provides a consistent installation across the facility. This is one of the most important items for the facility. Many installations are not provided the proper time for installation, configuration, testing and documentation. It is imperative that your deployment procedures address physical and logical items and that no device is installed "in a rush." It is also important to have standard operating procedures for decommissioning of devices as well. The testing and development networks should be complete separate from each other, as well as the production network. The testing network should emulate the production network as close as possible to ensure testing accuracy.

b) ***Monitoring*** – The ability to remotely identify issues within the data center, on a 24 x 7 x 365 basis is imperative to the operation of the organization. Monitoring touches all elements of the facility from the physical layer (power, cooling, temperature and humidity), through the network layer (routing, switching, carrier access) to the server and application levels. The ability to monitor all systems, receive various levels of alarms indicating degraded performance and/or outages and the standardized procedures on how to respond during these situations saves a vast amount of time when a trouble occurs. The organization should also provide for local and remote access to systems so that troubleshooting, analysis and corrective measures can be corrected remotely, when required. University of Utah should integrate Engineering and Operations best practices along with Network Operations and Monitoring Center procedures.

c) ***Troubleshooting*** – When an application, or device, experience degraded performance (or an outage) methodical troubleshooting and analysis is critical to the restoration of the service. Many outages have been addressed via the "chicken little" method where exorbitant amounts of energy, changes and measures are performed, sometimes without resolving the situation. This approach may breach security policies, actually reduce service or even eliminate availability. Having a structured and methodical approach that is taken during every occurrence will greatly improve the time to restore as well as the method and process in which service is restored.

d) ***Warranties and Service Contracts*** – Proper warranty and service contracts can save the entire facility under a major outage. The ability to not properly service items could lead to total facility outage and the inability of the organization to function from an information technology perspective, thus causing revenue loss and loss of other mission critical services. Having the proper services in place to address facility infrastructure devices, hardware failure, software updates, and proper maintenance of the facility itself, provides an extended life for the data center as well as the applications and platforms hosted within.

e) ***Maintenance*** – Maintenance of systems is critical to ensure proper functionality as well as providing extended life for all systems. Maintenance includes all layers of operations within the data center and should be addressed at each tier. All maintenance efforts should be performed on a cyclical basis, by technicians and engineers qualified on the specific system and documented via a maintenance log. Maintenance should be scheduled during a specific frame of time (including start and stop). All end users should be notified if any maintenance will impact services or operations during the maintenance window. Many maintenance windows are scheduled during off hours so that services and end users are not impacted during the normal business day.

f) ***Asset Management –*** IT asset management (ITAM) is the set of business practices that join financial, contractual and inventory functions to support life cycle management and strategic decision making for the IT environment. Assets include all elements of software and hardware that are found in the business environment. Software Asset Management (SAM) applies to the business practices specific to software management, including software license management, configuration management, standardization of images and compliance to regulatory and legal restrictions—such as copyright law, Sarbanes Oxley and other contractual compliance. Hardware asset management entails the management of the physical components of computers and computer networks, from acquisition through disposal. Common business practices include request and approval process, procurement management, life cycle management, redeployment and disposal management. Asset Management is critical to the financial component of a data center. Many items are depreciated within an organization over a period of time. The ability to identify a device and it historical use has many benefits in an operational environment. Organizations that have strong asset management procedures typically use a database and physically track items by utilizing wireless scanners for ease of inventory and locations. Radio Frequency identification (RFID) may become more of an option for Asset Management in the near future.

Additional Considerations and Comments

a) ***Hardware*** - The hardware element should be one of the primary areas of focus during the standardization effort due to the fact that this facet encompasses a large portion of the infrastructure. This piece should be applied to the server, network, storage and other hardware components. Ease of engineering, implementation, operations, maintenance and asset management will be obtained through standardization. This also makes procurement and warranty services easier and most likely savings can be reached to due consolidated service agreements. The organization should break the hardware down to common component types based on functionality and determine a standardized deployment mode for those relevant functions. Accommodate for the occasional “one off” scenarios and limit them as much as possible. Ensure that the chosen deployment model can accommodate current and future needs.

b) ***Software*** - Properly maintaining software versions for network devices alone can be nerve wracking. The amount of versions, features, software bugs tracked, and security notices can consume a fair amount of a network engineer’s time

just staying on top of the proper version to use. The same concerns are relevant to the Windows Operating system and various server platforms. This is especially true in virtualized environments. By deploying consistent hardware platforms throughout the facility the amount of time is reduced from an operational perspective. It is highly recommended that all software is made available via secure means on the network as well as via an alternative method (disk or flash).

c) ***Configuration Management*** - Configuration Management (CM) focuses on establishing and maintaining consistency of a product's performance and its functional and physical attributes with its requirements, design, and operational information throughout its life. For information assurance, CM can be defined as the management of security features and assurances through control of changes made to hardware, software, firmware, documentation, test, test fixtures, and test documentation throughout the life cycle of an information system. Configuration Management enables consistent engineering, deployment, operations and troubleshooting methodologies. Many organizations are utilizing software suites to control configuration management and automating data center management across lifecycles, so that the IT organization can deliver cost-efficient and compliant deployments in complex environments.

d) ***Change Management*** - The objective of Change Management in this context is to ensure that standardized methods and procedures are used for efficient and prompt handling of all changes to controlled IT infrastructure, in order to minimize the number and impact of any related incidents required for data center service. Changes in the IT infrastructure, no matter how minute, may cause problems, outages, and can impact organizational requirements or goals. Change Management can ensure standardized methods, processes and procedures are used for all changes, facilitate efficient and prompt handling of all changes, and maintain the proper balance between the need for change and the potential detrimental impact of changes. A proper Change Management procedure and methodology is critical in data center environments. Change Management is actually a component of the Information Technology Infrastructure Library (ITIL). ITIL defines the organizational structure and skill requirements of an information technology organization and a set of standard operational management procedures and practices to allow the organization to manage an IT operation and associated infrastructure. The operational procedures and practices are supplier independent and apply to all aspects within the IT Infrastructure. Any organization seeking to establish standardized operational procedures within a data center should investigate and implement an ITIL-based change program within their organization.

e) ***Documentation*** - During the design and engineering phases of a service, documentation is critical and a fair amount of time is devoted to accurately documenting the proposed or future solution. Unfortunately, once the service is deployed, that final approved drawings are rarely updated over the lifetime of the IT service. As soon as an outage occurs one of the first questions asked by the person or team troubleshooting the issue, is “Where is the documentation”? Accurate documentation can save time for many facets of the Data Center and IT operations staff and an effort should be made to consistently produce accurate documentation during all phases of an IT service. An organization should define documentation procedures, methods, outlines and templates to ensure consistency and accuracy across all documentation packages. This methodology should also include version control as well as archiving of older sets for future reference. Clear and concise University of Utah technical documentation can be quite vital when the need arises.

f) ***Management Software*** - Many tools and software packages are available to manage IT infrastructure, practices and procedures across the enterprise as well as for data centers. Unfortunately, there is not a holistic framework that integrates all of the packages to act as one common tool. A number of organizations have successfully integrated the complex features and functionalities into one overarching software suite, sometimes called a Manager of Managers (or "MOM"). The MOM integrates multiple disparate software packages into one unified functional solution. The MOM can integrate all aspects of management suites, from networking, servers and SAN's to the outage and trouble ticketing tool, to the ITIL change and configuration management packages. This approach, although quite complex, can provide a holistic end to end view of the data center operations and the services that it provides. Many Executives depend on visual "dashboards" to keep a remote eye on the actual service delivery for the organization. Finding the correct software for the job, and the ability to integrate all packages into one solution provides the organization with an end to end solution that can reap great benefits once it is properly deployed and integrated.

Given that hardware delivery may involve pallet jacks rolling concentrated loads, Staging areas are often non-raised floor space in close proximity to the raised floor ramp (unless the data center has a depressed raised floor). EYP MCF recommends that considering construction of a Staging Area.

g) ***Data Acquisition*** - Power monitoring can be a very useful tool to optimize cabinet loading, load distribution, and cooling system effectiveness. There are several ways to approach power monitoring. The preferred method is dependent on the IT organization's commitment to use the information on a proactive basis. A few points to consider are;

1. Implement Smart-PDUs in cabinets that can measure power consumption at each individual receptacle allowing load management at the device level. This can be useful information when planning the installation or deinstallation of equipment. To make this investment cost effective, the IT organization must commit to load management at the device level.
2. The implementation of Smart-PDUs requires compliance throughout the life of the data center. This means that the organization must commit to the purchase and installation of a Smart PDU with every rack purchase.
3. Smart-PDUs require network connections with dedicated aggregation switches to provide remote access to all. Here again is another commitment the IT organization must continue through the life of the data center.
4. An alternative is to install branch circuit monitoring at the panelboard level within room-PDUs and room-RDUs. The initial cost may be higher that Smart PDUs but it eliminates constant Smart PDU procurement and installation costs. The branch circuit monitoring remains a fixed asset no matter how often branch circuit and receptacles require replacement. Branch circuit monitoring only provides cabinet level load data. The IT organization must determine if this is sufficient for their needs.
5. Many hardware manufacturers are implementing environmental monitoring points (temperature, power consumption, etc.), device level monitoring will eventually become available through operating system options with the ability to

script data acquisition and aggregation as desired. This may preclude the necessity and expense for Smart PDUs making branch circuit monitoring a more cost effective solution.

h) ***Data Management*** - To maximize the investment of power monitoring equipment, the IT organization must commit to proactive use of the data. Part of the hardware deployment standard should dictate deciding device location based on cabinet load and the ability to cool that device. This means using the load data to determine which cabinet can accept the new device without exceeding its own load limit and ensuring load uniformity with surrounding cabinets.

Maintaining uniform loads in cabinets optimizes the cooling system decreasing the need for over provisioning on cooling for one or two high-density cabinets standing amid several low density cabinets. Such a condition occurs when arbitrarily locating devices in any opening cabinet space without consideration for its impact to surrounding loads.

Cabinet level load data is usually sufficient to achieve this objective with minimal extra effort. In most cases, only very dynamic hardware environments would benefit from load data acquisition at the device level.

i) ***Standard cabinets -*** Cabinet standardization will greatly simplify many aspects of data center planning and management. Given the advent of rack mount solutions, cabinet standardization is simple to establish. Making the initial investment to populate the data center with as many standard cabinets as possible will simplify installation planning and allow for rapid deployment of hardware when unexpected project arise. Installing cabinets and related power distribution in advance, will increase the initial data center build cost but will eliminate the need for short-term funding requests and limit the risk that funding will not be available. Not all equipment will fit the rack mount model. The University must leave some empty space for standalone type equipment cabinets.

For rack mount style cabinets, EYP MCF recommends the university consider the following attributes when selecting a standard cabinet.

- 28-inch wide racks to provide cable management space for extra long cable dress
- 42-inch depth to accommodate the increasing depth of servers and for unobstructed cable management directly at the rear of device without impacting airflow
- Purchase filler panels for all open RU spaces at the front of the cabinet
- Ensure there are no gaps between the front mounting rails and the sides of the cabinet where device exhaust at the rear can infiltrate the front of the cabinet. This condition is quite common and negates the value of filler panels.
- Do not purchase (cabinet) ceiling mounted flushing fans. These fans do little to extract heat from the cabinet. However, an entire row of cabinets with operational ceiling fans will create an air curtain between the top of the racks and the ceiling that will impede or redirect airflow returning to air-conditioners.

- Ensure cabinet doors offer a minimum of 63% open area (perforations) to ensure sufficient airflow
- Purchase options to block airflow between cabinets bolted together
- Purchase cabinets with top cable entry access
- Consider cable management options to promote quality cable dress

j) ***Standard Cabinet-PDUs -*** There are many types of Cabinet-PDUs (Power Distribution Units) available with increased capacity, flexibility, and functionality keeping pace with changes in hardware technology. Not long ago, 120-volt, single-phase circuits were sufficient. Today, multiple three-phase circuits are gaining rapid acceptance. As with cabinets, standardizing on PDUs will simplify power management moving forward. With the uncertainties of the future, selecting the most appropriate cabinet-PDU can be difficult.

The table below presents the typical power capacity available from the most common cabinet-PDU configurations available today. It also shows the quantity of each cabinet-PDU configuration required to satisfy typical cabinet loads. The quantities depend on input power redundancy. Some equipment may require redundancy and some may not. The table accommodates both conditions.

		Quantity of PDUs / Branch Circuits Required			
		Low Density Area – 8 kW/Cab		High Density Area – 12kW/Cab	
Cabinet-PDU Configuration	**Available Power (kW)**	**Non-Redundant**	**Redundant**	**Non-Redundant**	**Redundant**
120-volt, single-phase, 30-amp	2.8	3	6	5	10
208-volt, single-phase, 30-amp	4.9	2	4	3	6
208-volt, three-phase, 30-amp	8.6	1	2	2	4
208-volt, three-phase, 60-amp	17.3	1	2	1	2

Some attributes to investigate when selecting a PDU are;

- How many type of receptacles
- How many receptacles of each type
- Maximum load for a given segment of receptacles
- Maximum power capacity of the PDU
- How accurate are the current measurement, if power monitoring is available
- Are current measurements trueRMS
- Are current measurement per phase leg, per segment, or per receptacle
- Are circuit breakers clearly visible, easily accessible, and protected against accidental trip
- Does it have on/off remote control functionality; at the PDU, segment, or receptacle level
- Does it offer a simple network connection (RJ45, Ethernet) with web-interface
- How does it mount in the cabinet; U-space, zero U (side), vertical, etc?

Smart-PDUs offer a myriad of functionality but, as mentioned earlier, one must question whether the functionality is necessary or more importantly, will you use the features. Additional functionality increases the propensity for component failure.

EYP MCF recommends a simple vertical-mounted 208-volt, three-phase PDU with several C13 and at least three C19 receptacles and power monitoring with local display (likely to be a feature rather than option in the near future). Consider branch circuit monitoring at the room-PDU level as a cost-effective alternative to networking Smart-PDUs.

k) ***Raised Floor versus Non-Raised Floor -*** Few manufacturers design computer equipment that requires a raised floor for successful operation. This is especially true with rack mount style. Although raised floors may appear as an unnecessary expense, as with space they are not a major cost driver in the construction of a data center. Eliminating the raised floor cost will usually increase the cost of other systems required to replace the benefits of a raised floor, such as overhead power distribution and overhead cooling distribution. The cost of manual labor working overhead quickly offsets the savings eliminating the raised floor.

When comparing raised floors used for cooling or non-raised floors with overhead cooling, each has its cooling limitations. The best of either design begin to suffer at 8kW/cabinet (assuming uniform load for all cabinets). Higher cabinet densities usually require some form of supplemental cooling. Many of the supplemental cooling products available today do not require raised floors to perform their cooling, however the facilities (plumbing) to operate those units prefer raised floor. Most of the supplemental or alternative cooling systems available for high-density applications use water as the cooling medium. Locating water piping under a raised floor is preferable to overhead from an installation (overhead labor cost), accessibility (valves), and disaster (leak) perspective.

Raised floor environments offer greater flexibility for the unknowns such as;

- future computing hardware (standalone supercomputers) that might require under floor air or water cooling,
- future supplemental/alternative cooling equipment preferring under floor plumbing,
- future equipment with severe interconnect cable length restrictions (supercomputers with high-speed interconnects) where cables cab stretch farther under a raised floor versus overhead

7.3 Business Continuity Planning and Disaster Recovery Considerations

BCP Overview

Disasters can take many forms. While natural catastrophes like flooding, hurricanes or earthquakes may be infrequent events, more common causes of systems disasters such as system outages to computer viruses to disruption by discontented employees

can strike at any time. Resuming normal operations as quickly as possible minimizes business disruption, and good preparation will ensure that.

Many organizations and companies aren't adequately prepared for systems disasters. Recent research shows that major barriers to preparation include lack of Executive support and funding. Adequate funding for disaster recovery efforts requires a shift in priorities of an organization's IT initiatives. In the past, organizations implemented technology as a cost savings measure. Now, IT initiatives that support business continuity and revenue generation are getting top priority.

In the early days of data processing, the mainframe computer was usually housed in a large room with very large windows so everyone could see the computer. This led to the term "glass house." The term "Disaster Recovery" is usually related to only the restoration of the "glass house." In the same vein, the term "Disaster Recovery Plan" related more to a plan on how to restore the "glass house" and its contents in the event of a crisis.

In today's complex work environment, we not only have to take the concept of the "glass house" into consideration, but also the client/server computer networks and the work-areas where essential business functions occur. The work-area includes all the needed facilities such as desks, chairs, telephones, office supplies, and so on. Another often-overlooked aspect is the human factor. Any recovery efforts would surely fail without having an adequate number of trained personnel on hand to actually perform the critical business functions. Today's more encompassing recovery environment is usually referred to as "Business Continuity." A Business Continuity Plan (BCP) is defined as:

> A document containing the recovery timeline methodology, test-validated documentation, procedures, and action instructions developed specifically for use in restoring organization operations in the event of a declared disaster. To be effective, most Business Continuity Plans also require testing, skilled personnel, access to vital records, and alternate recovery resources including facilities.
>
> Properly written, a BCP is a collection of procedures and information which is developed, compiled and maintained in readiness for use in the event of an emergency or disaster. This would include the elements of a disaster recovery plan (DRP). Putting it simply, business continuity is the process of planning to ensure that an organization can survive an event that causes interruption to normal business processes.
>
> Disaster recovery is the process that takes place during and after an organizational crisis to minimize business interruption and return the establishment as quickly as possible to a pre-crisis state. The process of creating, testing, and maintaining an organization-wide plan to recover from any form of disaster is called Business Continuity Planning (BCP).
>
> Every BCP strategy includes three fundamental components: risk assessment, contingency planning, and the actual disaster recovery process. BCP should encompass every type of business interruption -- from the slightest two-second power outage or spike up to the worst possible natural disaster or terrorist attack.

The objective of disaster recovery planning is to enable an organization to recommence normal IT functions as quickly and as effectively as possible following a disaster or disruption to computing services. An impartial, bottom line assessment of the true impact of a systems disaster on an organization can quickly point out the need to be prepared. Up-front integration of BCP budgeting into all of the strategic University of Utah IT initiatives will help spread the financial overhead fairly among all users of IT systems. It will get people thinking about the importance of BCP as an essential ingredient of any computing initiative.

Creating and maintaining a workable business continuity plan (BCP) is an essential factor in ensuring the continued survival and prosperity of the University of Utah organization. It is highly recommended that University of Utah develop a full BCP and DRP for their entire enterprise.

Application Prioritization Ratings

Similar to the facility stratum, it is also recommended that University of Utah evaluate and rank all of their applications in a similar manner. An application prioritization system (i.e. high, medium and low) would provide a pre-defined set of services to support each specific application level.

All applications deemed critical (i.e. SAP, e-mail, applications with specific Sarbanes-Oxley requirements, etc…) for University of Utah operations would be rated at the highest level. These applications would be housed within the University of Utah data center(s) and would have a fully bevy of high availability services to support them (facilities, monitoring, full network services, options for blade and virtualization, storage, tape etc…). Further finite criteria specific to University of Utah operations and business objectives would establish what each rating consists of.

The prioritization of applications and services would also assist University of Utah in engineering, deployment, support, operations, troubleshooting and data replication as well as Business Continuity (BCP) and Disaster Recovery (DR) planning.

Attributes of enhanced BCP/HA/DR capabilities utilized in this study include:

Active/Active data center mirrors for both P1 and P2 rated application environments. The Primary site would host all production systems, but the HA centers site would host a sub set of defined critical systems.

Enhanced HA/DR site capabilities to include synchronous online data storage backups (SAN type technologies) for both P1 and P2 rated applications. The HA/DR capability of providing synchronous remote site data mirroring of critical production application data is quickly becoming a best practice for many industries.

<u>Assumptions of Attributes for Active/Active Data Center Mirrors:</u>

- Location of data center mirror within a 40km fiber distance of the primary production data center allowing current network technologies to provide synchronous application data mirroring.

- Optimized distance of the fiber connections between the primary and secondary mirror sites would be within 40km total distance. 100km is the maximum distance to ensure synchronous data mirrors according to Cisco Systems, but 40km is submitted as an optimum distance utilizing price/performance criteria.
- Best practices dictate diverse path dual network paths between the Active/Active data center sites.

As University of Utah plans the effort to consolidate data centers, applications, servers, and storage systems into consolidated data center facilities, it is critical to ensure that the Wide Area Network (WAN) bandwidth capabilities are maintained and/or enhanced for any data center consolidation efforts so that available data center services provide the highest levels of availability, reliability and scalability.

Network Based Operational Continuity Considerations

As an industry-standard design rule, the secondary data center facility would need to be within forty (40) kilometers of the primary production data center facility in order to satisfy synchronous network architecture requirements for successfully and consistently mirroring stored data between data centers with the highest levels of performance and reliability. This distance liability sometimes limits and organization from deploying a secondary data center or disaster recovery site.

There are various Cisco Systems solutions, such as DWDM and fiber-channel buffer-credit extension techniques, which claim to extend the distance between the primary production data center and the secondary facility without sacrificing synchronous performance.

HP's StorageWorks solution, using Data Replication Manager (DRM), can provide the capability to replicate data over direct fiber channel, covering distances of up to 100 km (~62 miles) via the Very Long Distance GBIC. With DRM, EYP MCF can replicate data at full fiber channel speeds. The use of an extended fabric license is recommended for additional buffer-to-buffer credits at these distances.

EMC reports that their Symmetrix Remote Data Facility Synchronous (SRDF/S) product ensures zero data exposure for remote data replication over distances up to 200 km. SRDF/S provides key functionality such as site failover and fail back, source/target dynamic switching, incremental restore with immediate access/updates and multi-patching support.

In addition to utilizing various technologies available from the various vendors, as well as the Network Carriers, University of Utah should also investigate the various data compression, load balancing and acceleration technologies that are available. These technologies can enhance data throughput and increase availability for data, SAN and other mission critical services. These appliances should be thoroughly tested with the existing technologies utilized within the University of Utah network, and can provide a noticeable improvement in performance as well as a real return on investment.

One solution that could assist in this manner would be the Riverbed Steelhead 6120. Depending on a number of factors, certain deployments with the Riverbed appliance can expect to see benefits in the range of 5 to 10 times faster performance over the

WAN. In addition, each 6120 unit can be expected to deliver unrestricted throughput on the LAN-side, and up to either 310 or 800 Mbps (high speed) on the WAN side.

Multiple Steelheads can be clustered and load balanced to achieve failover, redundancy, or simply higher overall throughput of the system. Using a product such as F5's BIG-IP Global/Local Traffic Manager (GTM/LTM), in conjunction with the Steelhead, could yield N+1 scalability and only is limited by the size of the WAN connection. This throughput could be scaled as large as an OC-192 (10 Gbps) circuit.

Utilization of these various products could enhance throughput, reduce latency, increase availability and reduce costs for the WAN. To utilize these solutions would require detailed testing evaluation, engineering and cooperation from the Network Carrier and Network Vendors. These solutions are also dependent on many physical layer aspects such as the type of buried fiber currently available, distance, quantity of cross connect points and quantity of splices. Needless to say, careful analysis of the end to end architecture would need to be performed to ensure the success of providing the high speed synchronous connectivity beyond 40 kilometers.

7.4 Data Center Facilities Considerations

Staging Area

Enterprise data centers typically dedicate a space, separate from the computer room proper, referred to as a Staging Area or Server Build Room. The intent of this space is to accept delivery of hardware, unpack and inspect the equipment, prepare it for deployment, and deliver to the data center. The most common benefits to such an area are;

- Prevent the dust, dirt, and debris associated with packaged hardware from contaminating the data center environment
- Minimize unnecessary traffic of server build personnel in the data center
- Minimize extended periods of time blocking airflow to critical equipment in high density areas with boxes set on vent tiles, doors open, etc.
- Minimize the risk of damage to operational hardware with build activity in the data center (i.e.; loose hardware falling into hardware)
- Provides isolated power in the event that new hardware arrives with electrical problems (not unheard of). Short circuits would upset Stage Area power but not data center
- Provides comfortable working area for server builders to load operating system and applications, test network connections and ensure server is ready for deployment into data center and immediate use.

These benefits demand that the power and cooling sources for the Staging Area be separate from those of the data center. Given the minimal number of devices operating simultaneously, the building cooling system is usually sufficient for this space. The power distribution system requires sufficient separation from the data center so over current or transient events will not interrupt data center operation. The power distribution system requires sufficient flexibility to accommodate many receptacle types.

Given that hardware delivery may involve pallet jacks rolling concentrated loads, Staging areas are often non-raised floor space in close proximity to the raised floor ramp (unless the data center has a depressed raised floor).

Storage

Storage Rooms are often forgotten during the design phase or eliminated during budget review, yet they are the most sought after space once settled into the data center. There is always a need for hardware and accessory storage for a data center. Typically, the perimeter of the data center becomes the default storage area. The problems with perimeter storage are;

- Cardboard boxes contain most stored items. Cardboard is one of the single greatest sources for introducing dust and dirt into the data center.
- Cardboard also increases the fire load (amount of combustible items in the data center likely to sustain a fire).
- Excess perimeter storage often impedes service access to support equipment such as air-conditioners or power distribution units. In some cases it can impede access to critical plumbing under a raised floor.

EYP MCF recommends a storage room dedicated to the data center and in close proximity.

Physical Security

In addition to continuous and uninterrupted utility support, the establishment of a secure facility is an essential goal. Entrance points should be restricted to two locations. Access afterhours should be provided by a phone connected to a security desk. Also provided in the Lobby:

- Security guard presence.
- Secure access via a mantrap to operations center and the equipment spaces.
- Direct access to public restrooms.
- Camera presence.

Mantraps would include installation of exterior hand-geometry readers combined with assigned access codes. Devices would limit passage to one individual at a time. The Loading Dock is also considered a mantrap area, and access thereto would be controlled by a central security monitor in the Lobby.

Loss Prevention/Risk Management

- Physical compartmentalization of the spaces in a data center should be provided as follows:
- Sub-division of rooms housing redundant utility systems should be accomplished using minimum 1-hour rated walls of masonry construction to provide blast resistance.
- Separation of utility portions of the facility from the data processing areas should be provided.

Building Envelope issues include the following:

- Glazing should be limited to the Lobby and office areas.
- A highly wind-resistant multi-ply roofing system should be employed over lightweight concrete topping.

- Roof penetrations should be minimized. Roof drainage should be designed in such a way as to minimize the potential for standing water. Roof structures should be designed to support the weight of standing water that could result due to the failure of a drain.
- Plumbing vents and mechanical exhausts should be through walls.

Fire Protection

A cross-zoned, double interlock, pre-action fire sprinkler system, should protect all critical areas. A high-sensitivity smoke detection system (HSSD) should be utilized to provide an early warning of potential fire/smoke events.

Piping and Drainage

Chilled water mains should be routed in trenches, and the trenches equipped with floor sinks acting as drains and acting as termination points for condensate drains. Trenches should have continuous line leak detection. Trench drains, dams, and other devices should be employed to contain battery acid spills, overflow from toilet rooms, and other fluid hazards.
Pressurized piping containing fluids should be avoided over access floor areas as well as over the electrical equipment serving it.

Environmental Issues

Environmentally responsible refrigerants should be employed for the chillers. Underground fuel tanks and tanks integral with emergency generators should be double-walled and provided with leak detection systems. Underground fuel piping should be double-walled and provided with leak detection and may be installed in a concrete tunnel.

Care shall be taken so that generator exhaust is not drawn into the building air intake systems and so that soot discharge is managed. Exhaust discharge routing must be considered in locating generators.

Acoustical Considerations

Considerations should be made with regard to the chillers and the emergency generators. Local jurisdictional limits specific to a particular site must also be considered.

Ambient noise within the equipment areas, especially from CRAC units, should be reviewed. The data processing areas should provide an acoustically comfortable working area for employees.

Employee Welfare Issues

Appropriate ventilation, introduction of outside air, and filtration of air should be provided in areas occupied by employees to maintain good indoor air quality. Unconditioned outside air, however, should not be introduced to data processing areas in order to maintain strict humidity control.

Lighting should be designed to minimize glare in work areas.
Shower and restroom facilities for employees should be provided separate from those provided for visitors. Lockers should be provided for employees.

Flexibility of Use

The configuration of data processing equipment areas should be generally open and column-free to allow for ease of equipment installation and flexibility of layout. In any case, column spacing should not be less than 30 feet on centers. The overhead clear height in the equipment areas should be a minimum of nine feet from the top of access flooring to the finished ceiling line.

LEED

LEED is a rating system that quantifies the energy efficiency and environmental principles of new and existing building design. There are nine different LEED rating systems for different types of buildings. They are LEED for New Construction, Commercial Interiors, Shell and Core, Existing Buildings, Homes, retail, school, healthcare and Neighborhood. This document will focus on LEED for New Construction (LEED-NC). LEED-NC organizes environmental strategies into six categories, Sustainable Sites, Water Efficiency, Energy & Atmosphere, Material & Resources, Indoor Environmental Quality, and Innovation in Design. Points are earned for meeting criteria and four levels of certification can be reached. A new building is Certified at 26 points, Silver at 33 points, Gold at 39 points and Platinum at 52 points out of a total of 69 points. Prerequisite points in each category must be achieved for LEED certification. LEED for New Construction is a method for architects, engineers, building owners and operators to implement environmental designs and operational policies to leave their building with a lower environmental impact. Strategies include developing programs to use alternative transportation, minimizing energy and water usage, using environmentally preferred refrigerants in cooling equipment, choosing building materials from local manufacturers, and maintaining ongoing indoor environmental quality. Please see the LEED for New Construction v2.2 Scorecard for a tabulated display of recommended credits to achieve at each LEED classification. The scorecard incorporates suggestion of various points for data center operation. This scorecard should be only used as a template for the project inception. Detail analysis of each credit should be performed with respect to individual project characteristic.

There are some LEED credits that would be especially easy for a campus data center to achieve. For instance, Development Density and Community Connectivity awards a point for being nearby existing buildings and community services such as stores, banks and libraries. By selecting low-flow plumbing fixtures and landscaping with plants that require little to no irrigation, water usage can be significantly reduced, which can earn both Water Use Reduction and Water Efficient Landscaping credits. The Recycled Content credit can be earned by selecting building materials with recycled content. Additionally Construction Waste Management requires diverting at least 50% of construction waste away from the landfill by recycling unused building materials, salvaging materials on site or donating materials to a charitable organization. Depending on availability, Material Reuse credit might be achieved by reusing 5% of building materials salvaged from another building. Many credits are available for indoor air quality and are good engineering practice for data centers. Outdoor Air Delivery Monitoring, Increased Ventilation, Construction IAQ Management Plans and Low VOC emitting materials are all practical points to achieve. Depending on the layout and design of the administration spaces, points can also be earned for natural day lighting and views. One point is automatically captured if a member of the project team is a LEED accredited professional. Innovation points are available for going above and beyond the requirements of the credits listed. Innovation credits can also be earned for established green building designs and practices not covered by the LEED credits. Ideas for these credits may include committing to environmentally friendly cleaning products and maintenance policies or using chemical free condensate water treatment.

Even though data centers have more electrical load than most building types they are certainly still capable of achieved LEED status. Care must be taken to assure the most practical and cost effective means to meet the requirements of the credits without sacrificing data center performance. Up to 10 LEED points are available for Optimizing Energy Performance. To earn these points the building must demonstrate a certain percentage of energy savings versus a baseline building. More points are earned as a higher percentage of savings is achieved. As of June 26th 2007, LEED – New Construction requires that a minimum of two points are earned under this credit, which corresponds to a 14% savings in energy of the proposed building versus an ASHRAE 90.1-2004 baseline building. There are a number of strategies to reduce the amount of energy use within a building including occupant behavior, building operations, high efficiency equipment, high efficiency lighting, site shading and high R-value building shell.

Choosing precisely which LEED points to capture will depend on site location, architectural and engineering building design and level of LEED rating desired. It is important to incorporate LEED objectives into the design early in the design process to ensure that high levels of detail and coordination are achieved as many credits impact several disciplines.

7.5 The Uptime Institute Classification (from the Uptime Institute)

The Uptime Institute (TUI) classifies the performance and reliability levels of data centers as Tier I to Tier IV as described below. The existing University of Utah production data centers vary throughout the campus. In moving forward with new production data centers, a standard of Tier III Concurrently Maintainable data centers are recommended. This is in accordance with University of Utah current initiatives, and the standard set with the recent Computer Building upgrade.

The four Tiers are defined as follows:

Tier I - Basic Non-Redundant Data Center

This tier of data center is not a continuously operating facility. This tier of data center is susceptible to disruptions from both planned and unplanned activity. A basic data center must be shut down completely on a regular basis to perform any maintenance and repair work. Urgent situations may require unscheduled shutdowns. Spontaneous failures of site infrastructure components or distribution paths will cause a data center disruption.

Tier II - Basic Redundant Data Center

This tier of data center is not a continuously operating facility. This tier of data center is also susceptible to disruptions from both planned and unplanned activity. Except for maintenance of UPS modules and other redundant capacity delivery components, a basic redundant data center must be shut down completely on a regular basis to perform maintenance and repair work to the distribution systems. Urgent situations may require unscheduled shutdowns. Spontaneous failures of site infrastructure distribution paths will cause a data center disruption. Unexpected failures of capacity components may cause a data center disruption.

Tier III - Concurrently Maintainable Data Center

This tier of data center is a continuously operating facility. This tier provides for any planned activities to be conducted without disrupting the computer hardware operation in any way. Planned activities include preventive and programmable maintenance, repair and replacement of components at end of their life, addition or removal of capacity components, testing of components and systems. Redundant components and alternate pathways allow maintenance of all systems and equipment, replacement of components and eliminate most single points of failure. This requires sufficient capacity to carry the full load on one path while performing maintenance or testing on the other path. Spontaneous failures of facility infrastructure distribution paths will cause a data center disruption.

Tier IV - Fault Tolerant Data Center

This tier of data center is a continuously operating facility. This tier provides the ability of the site infrastructure to sustain at least one unplanned failure with no critical load impact. A Tier IV facility requires two active power and cooling paths. The two power paths need to extend to the dual cord IT equipment level. Static transfer switches are theoretically not required in a Tier IV facility but are generally provided for operational purposes. Fault tolerant functionality also provides the capability to permit any planned activity to be conducted without disrupting the critical load in any way. Any component is able to fail without disruption to the load.

7.6 Data Center Concurrently Maintainable Redundancy Attributes

A concurrently maintainable data center has redundant capacity components and multiple distribution paths serving the site's computer equipment. Generally, only one distribution path serves the computer equipment at any time. Each and every capacity component and element of the distribution paths can be removed from service on a planned basis without causing any of the computer equipment to be shut down.

The operational impact:

- The site is susceptible to disruption from unplanned activities.
- Planned site infrastructure maintenance can be performed by using the redundant capacity components and distribution paths to safely work on the remaining equipment.
- In order to establish concurrent maintainability of the critical power distribution system between the UPS and the computer equipment, Tier III sites require all computer hardware have dual power inputs as defined by the Institute's Fault Tolerant Power Compliance
- Devices such as point-of-use switches must be incorporated for computer equipment that does not meet this specification.
- Operation errors or spontaneous failures of site infrastructure components may cause a data center disruption.
- During maintenance activities, the risk of disruption may be elevated.

7.7 Data Center Maintenance Assurance

Proper maintenance of all data center equipment is vital to insure that operational model uptimes are maintained. It has been found in most cases that proper equipment maintenance and procedures can actually increase a data center's tier level. Listed below are operational reliability requirements that should be performed at a minimum to insure uptimes for mission critical facilities.

a. Survey the organizational structure annually to determine if the staffing and management resources assigned to the data center are sufficient and appropriate to achieve the desired level of availability.
b. Review the personnel job descriptions and evaluations to determine if the proper skill sets and competencies are clearly identified.
c. Survey the personnel evaluation data to determine if the staff competencies and skills align with those identified as necessary to meet the operational requirements of the site.
d. Review the training system employed to develop skills and competencies of the staff to determine the effectiveness of the program in maintaining an optimum operational performance level staff.
e. Review the site documentation annually to determine if it includes appropriate components for Standard Operational Procedures (SOPs), Methods of Operations Procedures (MOPs), Emergency Response Procedures (ERPs), and programmed alarm responses.
f. Review the historical site documentation annually to determine if it is complete with regards to original as-built designs, Original Equipment Manuals (OEMs) and commissioning records which form the essential knowledge base on how the data centers are intended to perform, be maintained and operated.
g. Review the Change Control Process as it relates to the management of the infrastructure systems, including task descriptions, identified risks, risk mitigation plans, inclusion of MOPs, SOPs, and ERPs where necessary to define tasks, and the manner in which vendor supplied labor is managed and validated.
h. Review Maintenance Management Practices to determine:
    i. If appropriate levels of spare parts are maintained.
    ii. If the Computerized Maintenance Management System (CMMS) is fully utilized for optimized equipment performance.
    iii. If Preventive Maintenance practices reflect current industry best practice strategies for respective data center Tiers.
    iv. If the Building Management and Control System are being fully utilized to determine maintenance reliability trends and threshold performance levels which would trigger remedial maintenance ahead of impending failures.

v. Survey the site annually to identify any “lurking vulnerabilities” which could lead to downtime based on human error, such as EPO (Emergency Power Off), Switch labeling, incorrect panel schedules, unlocked panels, dead wire blocking under floor cooling, etc.

From the surveys and reviews conducted above, analyze the ability of the staff to achieve or surpass the expected availability based on our experience in data center design, operations, and assessment data from similar sites. This should be compared to “best-in-class” operational procedures and practices.

From the survey, reviews and analysis of the data, a basis for operational improvements can be formulated which will assist the staff in achieving availability goals. These operational improvements fall into three categories:

1. Items which can and should be addressed immediately to reduce identified urgent risks,
2. Items which will require moderate investment in cost and time but which will contribute substantially to mitigating identified operational risks.
3. Items which will require substantial investment in staff, contracts, and procedural development to completely address issues identified in the analysis of the survey data.

EYP MCF has a division called Critical Facilities Assurance (CFA) which is a team of engineers who specialize in operational improvements through organizational analysis and maintenance procedures. The CFA division provides services to verify, support and enhance the inherent reliability of a mission critical facility design through commissioning, risk management, maintenance design, and testing services.

7.8 Power Usage Effectiveness (PUE) and Data Center infrastructure Effectiveness (DCiE)

At its core, PUE and DCIE are mainly facility-based performance metrics. The primary goal in determining the efficiency of a data center revolves around determining the efficiency of the cooling and power distribution systems (a.k.a. "overhead" systems). The good news here is that there are many recognized building efficiency standards that are available that outline a process for establishing building energy use performance. Even better is these standards also present approaches on how to compare buildings to a minimum energy performance metric (not unlike PUE and DCIE) even if there are differences in climate and HVAC system type. Also, these standards are often already a part of many building codes that are adopted by municipalities. So why re-invent the wheel? Any standard developed for PUE and DCIE needs to be built on a foundation coming from already-established energy standards.

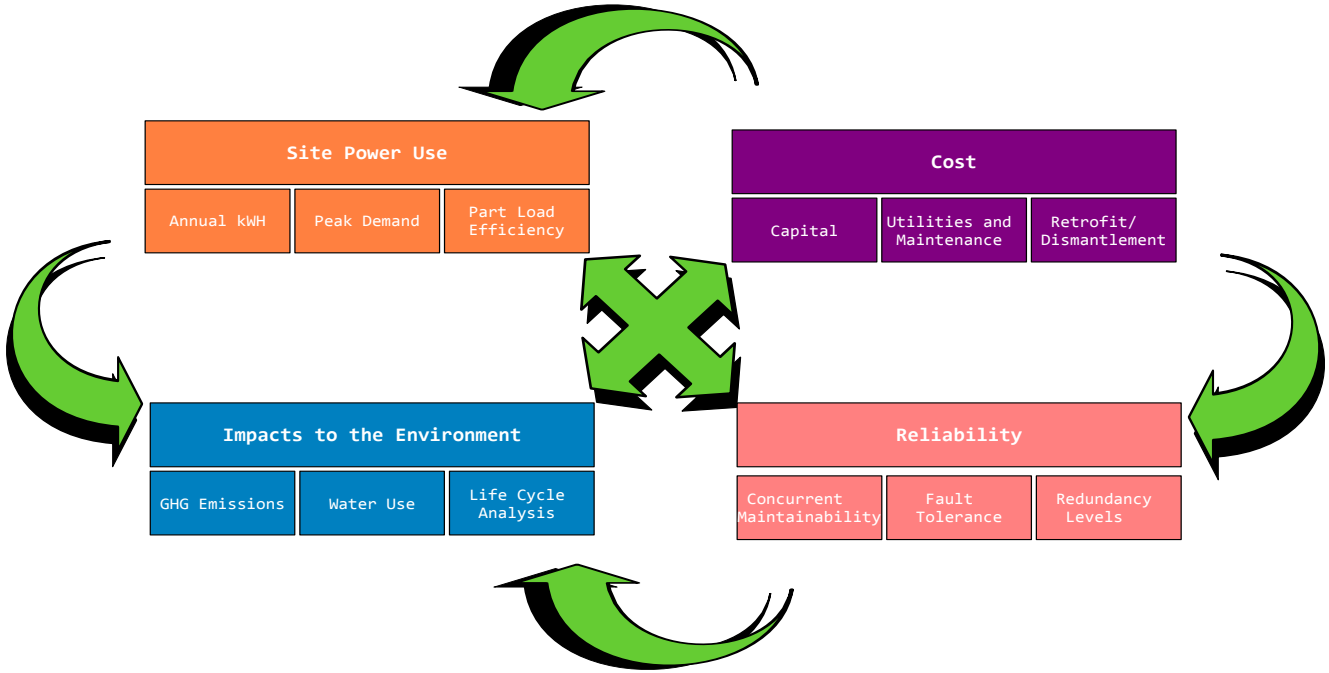

Figure 7: Understanding the interdependencies unique to data center planning, design and operations

As an example, using the methodology outlined in the ASHRAE 90.1-2007 energy standard, the primary components in a data center facility (or any commercial office building) that are analyzed to determine overall building energy performance are as follows:

Building Envelope

Although relative to the energy required to cool and power the ICT equipment, the energy impact from building envelope is small. However, basic code compliance and understanding the effects of moisture migration cannot be overlooked. For data centers, the integrity of the building's vapor barrier is extremely important as it safeguards against air leakage caused by the forces of wind and differential air pressure. It also minimizes migration of moisture driven by differential vapor pressure. Most data center cooling equipment is designed for sensible cooling only (no moisture removal from the air). Higher-than-expected moisture levels in the data center will result in greater energy consumption and possible operational problems caused by excessive condensate forming on the cooling coils in the air handling equipment. The energy strategies for the building envelope are specific to the climate and are presented in detail in the ASHRAE and CIBSE standards.

HVAC, Lighting and Power Systems

The energy standards present very specific requirements on the energy performance of HVAC and lighting systems, but very little on the power distribution systems as they are applied to a data center facility. Not having developed standards on UPS and the overall power delivery chain efficiency (from incoming utility power right to the individual piece of ICT equipment) is a major gap that needs to be filled. These standards do address, in great detail, how to judge the energy performance of HVAC and lighting systems, including control strategies, economizer options and climate-specific topics.

Focusing on HVAC, the biggest non-ICT energy consumer in a data center facility, the standards present minimum energy performance of individual components such as chillers, DX systems, pumps, fans, motors and heat rejection equipment. In order to be in compliance with the standard, it is mandatory that the equipment used meets the specified energy use metrics.

Since the largest power consumer in the mechanical system is the chiller (or other type of heat rejection equipment), one primary strategy to decrease overall energy consumption is to elevate the supply air temperature by increasing the chilled water supply temperature and/or reducing the temperature of the air moving across the condensing coil. However, the ability to incorporate this strategy will completely depend on the type of mechanical system, the climate and the allowable supply air temperature for the IT equipment. Consider that for fixed speed chillers, every 1 deg F increase in chilled water temperature can increase chiller energy efficiency 1-2 percent. For VSD chillers, every 1 deg F increase in chill water temperature can result in a 2-4 percent efficiency increase. Therefore, increasing the supply air temperature from 60 deg F to 75 deg F will result in an average efficiency increase of the chiller of nearly 40%.

Table 1: ASHRAE 90.1-2007 Allowable Chilled and Condenser Water Pump Power

Pumping Equipment Type	ASHRAE Allowable GPM/ton (GPM/ton)
Chilled Water	2.4
Condenser Water	3.0

Table 2: ASHRAE 90.1-2007 Allowable Fan Power

Baseline Fan Motor Brake Horsepower		
Supply Air Volume	Constant Volume (Systems 1 – 4)	Variable Volume (Systems 5 – 8)
<20,000 cfm	17.25 + (cfm - 20000) x 0.0008625	24 + (cfm - 20000) × 0.0012
≥20,000 cfm	17.25 + (cfm - 20000) x 0.000825	24 + (cfm - 20000) × 0.001125

Table 3: ASHRAE 90.1-2007 Allowable Power for Heat Rejection Power

Heat Rejection Equipment	Minimum GPM/HP
Propeller or Axial Fan Cooling Towers	**38.2**
Centrifugal Fan Cooling Towers	**20.0**

To determine whole-building energy performance, the ASHRAE standard sets forth a procedure that prescriptively defines how a given building's energy performance (the "proposed" building) compares to the calculated theoretical energy performances (the "budget" building). Indirectly, this method is the one used for the Energy and Atmosphere credit category in the LEED rating system, so as more data centers look to become LEED certified, this process needs to be used anyway. This same method (with some augmentation to address data center specific design and operations issues) can and should be used in determining the budget PUE and DCIE to benchmark data center energy use.

PUE Estimates for Data Center Facilities

Understanding and being able to analyze the primary energy consumers within a data center facility is crucial when targeting energy reduction/optimization strategies. These consumers come in the form of cooling and power distribution systems. Metrics for power use include power usage effectiveness (PUE) and data center infrastructure effectiveness (DCiE).

It is important to note that climate, cooling system type, power distribution topology and redundancy level (reliability, availability) will heavily influence the power use efficiency of the cooling and power distribution systems. These metrics are represented by the equations:

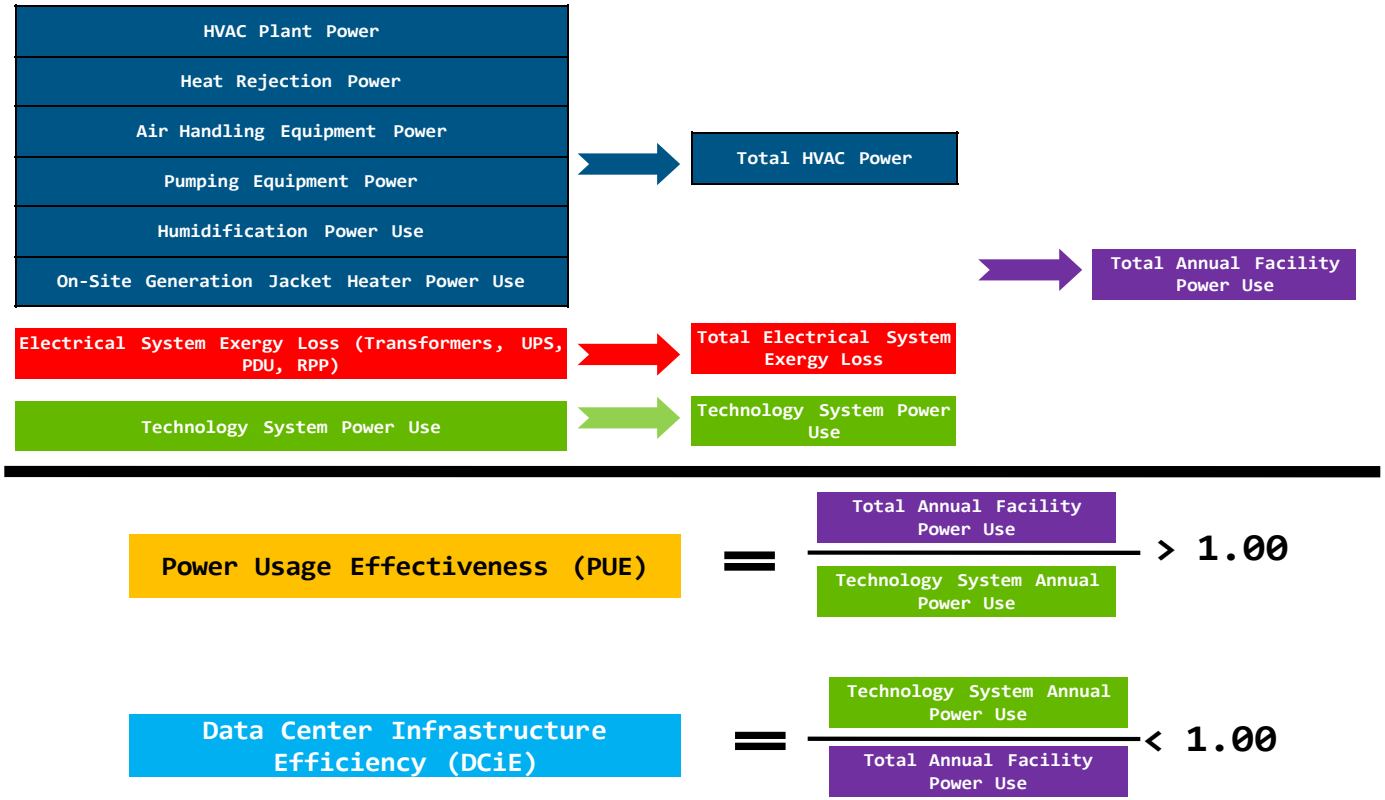

Figure 8: Components of PUE and DCIE

The following figures indicate estimated monthly PUE values and estimated equipment purchase and operating costs. These costs are **NOT** representative of the Program Recommendation and are presented for comparative purposes only for evaluating benefits of PUE. The comparison below is based on: 1,900 kW critical IT load x 110% = 2,090 kW total cooling load = 594 tons (Salt Lake City elevation).

Cooling Options Cost Comparison

Cooling Option	System Description	Initial Purchase Cost	Total Load	PUE	Annual Operating Cost ($0.10/kWh)	Return on Investment (compared to Option 1)
1	Conventional chilled water system without economizers. Air delivery by raised floor mounted CRAH units and liquid cooling systems.	$8,382,500	3,318 kW	1.75	$2,906,570	N/A
2	Conventional chilled water system with waterside economizers. Air delivery by raised floor mounted CRAH units and liquid cooling systems	$8,638,250	3,122 kW	1.64	$2,734,870	18 months
3	Conventional chilled water system. Air delivery by rooftop AHUs with outdoor air economizers	$10,325,000	3,089 kW	1.63	$2,705,960	9 years 8 months
4	Conventional chilled water system. Air delivery by rooftop AHUs with evaporative cooling systems.	$11,225,000	2,661 kW	1.4	$2,331,040	5 years

PUE and Costs are ROM estimates (±30%) based on the proposed systems and are for reference / comparative purposes only. Detailed analysis would be required to discover actual PUE and costs.

Note that due to the Tier 3 requirement, the chiller plant is retained at full size regardless of the outdoor air economizer and evaporative cooling systems. A system with a smaller chilled water plant would be less sustainable but would save considerable initial cost.

7.9 High Availability Storage Replication

RPO equal to 0

When an application is sufficiently critical that the business cannot afford the loss of any transactions, its RPO is defined as being 0. If an RPO of 0 is required, synchronous replication technology is the only solution available that can deliver it. Synchronous replication does not allow the PDC to process a transaction until the secondary data center signals that the previous transaction has been received. If the communication link between primary and secondary data centers is broken, synchronous replication can be set up to prevent the PDC from processing transactions until the link is repaired or replication is disabled.

The performance of synchronous solutions is sensitive to the network latency between primary and secondary data centers. The longer the distance, the more latency is induced in each transaction. For this reason, EYP MCF recommends synchronous replication only for distances up to a maximum of 40 kilometers.

RPO greater than 0

An RPO greater than 0 indicates that the administrator is willing to sacrifice some amount of data in the event of a disaster due to the use of asynchronous replication. For this type of solution, the asynchronous or journal replication options are best. Although complete data currency is not assured for this solution, data consistency may be provided when using enterprise class network based storage solutions from major solution providers.

Recovery Time Objective

A recovery time objective (RTO) indicates how much time will pass before applications are available for users again. Any disaster recovery solution requires a defined RTO.

RTO - Between One and Five Minutes

This setting indicates that applications on the disaster recovery data center systems must be automated with very fast reaction time to any disaster, not waiting for in-flight data to arrive. This requirement normally indicates cluster integration.

RTO of One Hour or More

This setting indicates that the administrator is willing to first assess the disaster and any possible data loss before initiating a data center failover and application recovery. Recovery can be automatic, but in most cases, a push-button (single command failover) approach is used, after the decision to recover is made. Part of the reason behind such a large RTO may be the time necessary to enhance the recovery point before starting the application in the recovery data center.

With more research, EYP MCF and EYP MCF can assist University of Utah in further defining long term, highly available Storage solutions (synchronous or asynchronous) to meet their business needs. The aforementioned storage solutions could be ideal for the University of Utah future state environment based on various data center strategy discussions and approaches that were discussed during this engagement.

7.10 High Density Computing

The latest trend of blade chassis server systems is increasing the per cabinet power loads to values in the 10 KW to 20 KW per server cabinet range. 20 KW per server cabinet power loads in excess of 500 Watts per square foot of raised floor space (localized power utilization) are possible. Legacy data centers often are limited to 50 to 100 Watts per square foot or less of power supply capacity.

Great care and planning must be provided in implementing the latest highest density computing platforms. High density computing can be utilized for High Performance Computing and typical enterprise platform computing. High density computing areas should be designated and appropriate design efforts must be undertaken to support these high density areas. Additional power and cooling infrastructure will be required for designated high density computing areas to support the high power and heating loads high density systems introduce into the data center environment. In-rack or above-rack cooling will be necessary within the designated high density computing area to supplement the overall data center cooling system.

EYP MCF has found that most legacy data centers struggle to provide increased critical power and cooling capabilities due to the increasing power and cooling loads introduced by more dense/compact technology platforms. The rate of increase in watts per square foot and the associated increase in heat loads per square foot has been steadily rising over the past four years due to the significant power density requirement increases in modern computing platforms. EYP MCF does not see this trend of ever increasing power per square foot slowing down for the next three to five years. Computer manufacturers are working to reduce the per processor core power requirement, but at the same time these same manufacturers are placing more computing cores on each processor chip. The net effect of this increase in processor cores per processor chip is a continuation of higher power requirements per square foot of raised floor data center area.

7.11 Data Center Migration Strategy

Migration planning was not part of the scope of this project, however early involvement of the migration planning effort for any mission critical data center is a recommended best practice. Migration Planning is a critical piece of the data center program that must be considered very early in the overall planning process. The migration strategy outlined below is the same for any of the EYP MCF options.

In the past, migration plans were started once construction efforts were under way with not much time to adequately complete the required planning, testing and relocation efforts. However, with today's high-density IT/Network deployments and production applications requiring continuous operations with little tolerance for risks and downtime, migration strategies/plans should be at the forefront of critical data center projects with planning beginning and running simultaneously with programming, planning, design, construction and commissioning efforts. The following information represents a methodological approach that must be a part of the overall data center program in order to ensure success throughout all segments of the project and to ensure, in this scenario that of all the devices, applications, and systems from each site are migrated with no planned downtime.

Approach

The migration methodology is comprised of six (6) steps in order to provide the required structure and management to the overall migration strategy from planning through execution.

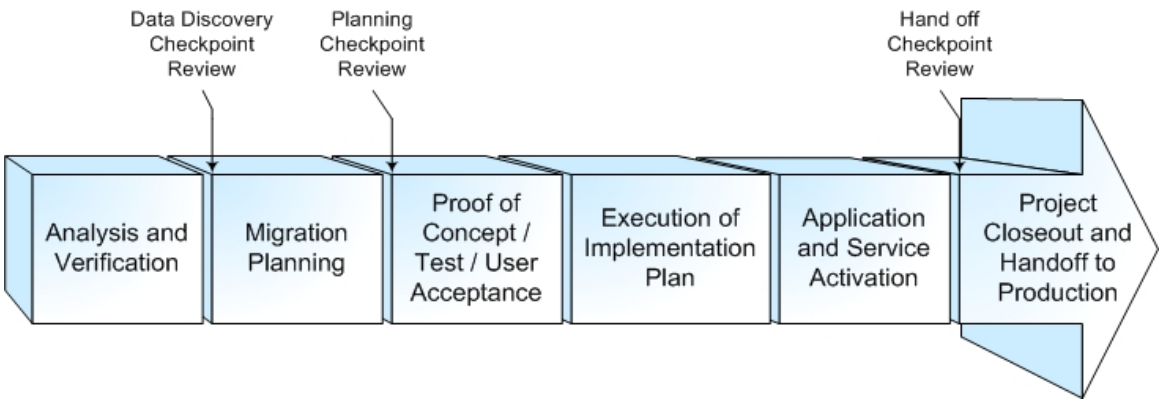

Figure 9: The Six Step Migration Process

A brief description of each of the Six Migration Steps is as follows:

Step 1: System Analysis & Verification

Step-1 will provide, through analysis, the information needed to thoroughly understand and quantify the current facility architectures and technology environments. Activities will include development of detailed project data, such as strategies, approaches, timing, resource requirements, scheduling, affinities/dependencies, costs, systems/applications inventories, and networks. The delivered output of this phase will allow management to make effective and final decisions regarding the migration strategy, approach, scheduling, and financial commitments.

Step 2: Migration Planning

Step-2 tackles the detailed project planning and development of the comprehensive plan strategy required for the actual project. This step affords an opportunity for University of Utah to review current configurations, operating models, operations and operational support methodologies and systems. This phase also allows for the development for all cut-over, back-out, contingency/Disaster Recovery Plans (DRP), a Business Continuity Plan (BCP), Business Impact Analysis (BIA), and Certification/Commissioning/Integration testing Plans and strategies. The activities within this step can be completed well in advance of the actual implementation of the Project Plan that occurs in Steps 3-5. A summary report is typically issued at the end of Step-2, which will reveal to all team members and executive management all of the required strategies, migration tasks, milestones, integration testing plans, and cut-over plans required throughout the overall migration planning phases.

Step 3: Proof of Concept/Unit Test/End User Acceptance

At the start of this Step, all physical and logical technology components will have been installed into the new facility and all local acceptance tests completed. This step represents a real world systems and applications environment within the new facility that can be accessed by a select group of users that are knowledgeable about how the real-world technology environment should function within full production mode. Integration Testing activities, such as inter-application access and updates, trial batch-processing jobs and system diagnostics under real world conditions, continue to be tested to simulate as closely as possible how the production environments scheduled for use will perform, function and react to the technology infrastructure, architecture and operational environment. Maintainability, operational support systems, and access-points are retested at this time to ensure near-real world conditions, as well as the ability for the operational organizations to perform all their required monitoring, maintenance, reporting and restoration functions. Upon completion of this Step, the overall production system/s will be certified as being ready to shift into a “go” or “no-go” decision completed by program and/or executive management.

Step 4: Execution & Implementation

Step-3 begins the actual execution of the planned activities which are the deliverables for Step-1 and Step-2. At this point all approvals (AFEs) have been approved and executed so that the actual project activities can commence. Step-3 can include tasks such as: purchase and installation of equipment; network implementation, systems software implementation; installation and configuring of storage systems and management; Command Center and/or Operations implementation; and Operational practices and procedures. Local acceptance and integration testing, as well as cut-over testing and planning, are key components within this overall step and must be completely developed and/or completed prior to the rest of the pending migration tasks.

A detailed Project Plan will depict all planned tasks, milestones, timeframes, start/finish dates, integration and proof-of-concept testing, cut-over plans, and team/staff resources for all relevant sub-projects. The Integration testing plan will reveal all testing processes, plans and methodologies for all IT/Network systems and applications, with all of these areas being also directed to all relevant business units for their specific testing and/or Proof-of-concept (POC) process requirements. The Cut-Over plan and strategy will depict minute-by-minute tasks required for each identified IT/Network "group" and/or sub-project. During this important step, it is very important to implement a management "dashboard" that can provide constant views as to the overall status of the migration tasks and milestones as the project moves forward.

Step 5: Application and Service Activation

Activation of all QA/TEST/DEV systems, applications and/or environments is the final execution Step within the Migration Planning Project Plan. This occurs after the proof of concept and integration and acceptance testing has been completed and the platforms and all associated production environments have been established. The final stages of the cut-over plan are executed so that all QA/TEST/DEV systems and applications can be made "live" within the new technology environment.

Step 6: Project Closeout

Upon the successful completion of the entire Migration Project plan and the final sign-off into the QA/TEST/DEV environment, this Step is the final migration component utilized to execute all remaining project activities, such as: disposing of obsolete equipment and/or facilities, verifying all punch-list and inventories, reassignment of staff individuals who were assigned to the Project Team, and officially closing the project within the Program and Project Offices.

Each Step will conclude with three common activities:

- Conduct a "Lessons Learned" session to gather and document "the Good, the Bad and the Ugly" - that is, what went right, what went wrong, and what still needs improvement.
- Recommendations for process and organizational improvements will also be gathered and, if approved, will be incorporated into the project plan for the next phase.
- If required, prepare a budget for the next phase for financial and personnel resources.

Conclusion

Many elements within the data center migration process must be properly planned for and maintained. It is critical that University of Utah consider these elements and integrate them as needed into current operational methodologies. By providing a structured and defined methodology, University of Utah will reap multiple benefits which will increase their overall facility availability, decrease operational costs as well as enable efficient operation.

8 Appendix A – Glossary

Chiller	A piece of equipment for removing heat from a gas or liquid stream for air conditioning and cooling purposes
Cold Site	A site in which space or equipment is available when needed
Computer Room Air Conditioning (CRAC)	A device that monitors and maintains the temperature, air distribution and humidity in a network room or data center
Computer Room Air Handler (CRAH)	A device that monitors and maintains the temperature, air distribution and humidity in a network room or data center
Consolidation	Initiatives to remove redundant hardware, software, maintenance and service costs out of data centers and thereby reducing the size and/or number of data centers needed
Direct Expansion Cooling (DX)	A cooling system that utilizes a refrigerant - like Freon - for cooling and dehumidification
Emergency Power Off (EPO)	Provides a single point of emergency equipment shutdown
Generator	A utility device that converts mechanical energy into electrical energy, available either in the form of direct or alternating current
High Performance Computing (HPC)	Compute intensive systems that require power densities of at least 10 Kilowatts per computer cabinet. Many HPC systems require greater than 20 Kilowatts of power per cabinet.

High Availability (HA)	The ability of a system to stay operational when hardware fails – allowing the compute load to be redirected to another system. This fail-over would not be noticed by the end user
Hot Standby Site	A site which mirrors the organization's production databases in real time
kilo-Volt-Ampere (kVA)	1000 Volt Amperes; a measure of apparent power
kilo-Watt (kW)	1000 Watts; a measure of real power.
Power Distribution Units (PDU)	An electrical device is used to control the distribution of power to the individual loads.
Rough Order of Magnitude (ROM)	Estimates within +/- 20% the cost of constructing the mechanical/electrical data center infrastructure
Signal Reference Grid (SRG)	A network of copper wires typically installed below a raised floor in a data center
Single Point of Failure (SPoF)	Any component that can cause a loss of critical load when it fails, and for which a countermeasure has not been implemented
Uninterruptible Power Supply (UPS)	A device that provides power backup when utility power fails or drops to an unacceptable voltage level
Virtualization	The creation of a virtual (rather than actual) version of something, such as an operating system, a server, a storage device or network resources.

9 Appendix B – Data Center Workshop Planning Notes

Project: **Data Center Improvements - Programming**
Client: University of Utah - Project 20109
Location: Eccles Broadcast Center, CU Conference Room
Purpose: Project Kickoff, Meeting No. 1
Date: May 12 to May 15, 2009

University of Utah Kickoff Meeting - 5/12/09			
Name	**Organization**	**Phone**	**Email**
[BB] Bill Billingsley	CDC	801.585.0073	bill.billingsley@fm.utah.edu
Joe Breen	Center for HPC	801.550.9172	joe.breen@chpc.utah.edu
Steve Corbato	CI Strategy	801.585.9464	steve.corbato@utah.edu
Jim Turnbull	CIO Hospitals and Clinics	801.585.7530	jim.turnbull@hsc.utah.edu
Glen Cameron	Data Center Manager	801.580.9920	gcameron@acs.utah.edu
Kenning Arlitsch	Marriott Library	801.585.3721	kenning.arlitsch@utah.edu
Mike Babinger	MLIB	801.581.8001	mike.basinger@utah.edu
[EL] Earl Lewis	Office of IT - PM	801.581.3635	earl.lewis@utah.edu
Dave Huth	OIT	801.585.9467	dave.huth@utah.edu
Andrew Reich	OIT	801.587.0902	andrew.reich@utah.edu
Jim Livingston	OIT/ITS	801.587.6085	jim.livingston@utah.edu
Brent Elieson	OIT/ITS	801.587.1320	brent.elieson@utah.edu
Caprice Post	OIT/ITS	801.585.5404	caprice.post@utah.edu
Stephen Hess	U of U CIO	801.581.6180	stephen.hess@utah.edu
Bryan Peterson	UEN	801.585.7789	bryan@uen.org
Gabriel Betit	Skanska	801.367.7925	gabriel.betit@skanska.com
Keith Hoover	Skanska	801.260.4660	keith.hoover@skanska.com
Doug Demmel	HP - Sales	303.604.6230	demmel@hp.com
Leland Gibbs	HP Services - Sales	602.549.8622	leland.gibbs@hp.com
[CP] Charles Prawdzik Jr.	HP/EYP - Architect - PM	310.689.3522	cprawdzik@hp.com
Peter F. Gmiter	HP/EYP - Architectural	310.689.3520	gmiter@hp.com
Sonny Siu	HP/EYP - Electrical Engineer	415.748.0508	sonny.siu@hp.com
John Tidd	HP/EYP - Mech Engineer	415.748.0503	jtidd@hp.com
Scot Hewth	HP/EYP - Sales	970.227.0869	scot@hp.com
[SC] Steve Carter	HP/EYP - Tech. Cons. - PM	312.343.9535	scarter@hp.com
Rob Myers	HP/EYP - Tech. Cons	312.909.1567	rmyers@hp.com

Italics indicates new discussion or corrections/updates/results/response from a prior meeting

Item	Discussion	Action/Date
1.00	**General and Administrative**	
1.01	Introductions	See Above
1.02	Project Kickoff Checklist	
1.03	HP [CP] Document Request Checklist – Not Reviewed → CP to provide information, followup with MCI + Med Records Designers → Contact: Mike Morgan (UU site purchase realator)	Info CP
1.04	HP [CP] Project Schedule Review	Info

Item	Discussion	Action/Date
1.05	HP [CP] Deliverables List to be posted	CP
1.06	HP [SC] High Level IT – DC Cost – Presentation → Looking at the economics of an evolving plan → Looking for Obama stimulus money as one source of funding for renovation and renewal of infrastructure → [UU] Operating costs of 7 to be (D) decommissioned sites may encourage their closure in order to justify cost of (N) Temple DC	Info
2.00	**Design Considerations**	
2.01	UU proposed destination Business (BU) entities: → CORE = ACS / OIT / CORE (UEN) / Hospitals → HDC = High Density Computing → COLO = Colocation ACS → Payroll, email, credit card, clinical → Must be 24/7/365 with ZERO downtime, concurrently maintainable → DR site in Richfield, UT (appox 131 mi south/southwest of UU) - Different seismic "zones" reportedly - Too great a distance for "Active/Active" application considerations UEN = UU Network Connectivity Hub; currently for nearly all UU colleges → UU seeks to provide added value of shared services with (N) Temple → Must be 24/7/365 with ZERO downtime, concurrently maintainable Hospitals → Must be 24/7/365 with ZERO downtime, concurrently maintainable → TBD on dedicated admin/office at (N) Temple ~ currently run by ACS; possible 7-8 staff required to migrate systems - Would not separate hospital from UU systems. Reliability requirements for both systems becoming more aligned → Daybreak in development stages: couple of clinics, ambulatory care - long term plan to evolve into inpatient hospital - Seeks consolidation and standardization over long term at (N) Temple → Currently 3 tiered levels of storage, considering a 4th - No classified research or SCIF requirement, although HIPPA likely - Shall contain huge-storage imaging files (Tera- and Petabyte size) - ±300 (E) applications running on Hospital systems - Network, wireless on Cisco gear currently COLO → Possible Tenants: EMS, ARUB (Med clinical lab), state/local & ed. orgs - UU has good relationship with State Officials	Info

Item	Discussion	Action/Date
	- UU understands Tenants to be "related" to UU - UU does NOT intend to compete with private, local Colo businesses - ARUP = 10th largest clinical lab has data center on campus; expressed desire to utilized (N) Temple DC → Could be used to pay for/offset (N) Facility costs - UU seeks effective cost modeling for rental space → Tier TBD, likely Tier III or greater HDC → Must be 24/7/365 with ZERO downtime, concurrently maintainable → Looking for 3 phase power systems with water to rack cooling CORE → Must be 24/7/365 with ZERO downtime, concurrently maintainable → Tier III MEP DCs → Monitoring of equipment required (M/E + IT) - CHPC is currently using monitored power strips and is a new standard for the university → Goal: never bring down an application for facility maintenance or outage → No requirement to separate physical network from rest of core; COLO TBD → ACS, OIT, Hospital Services (ITS) and CHPC are most important organizations for Phase 1 of (N) Temple → (E) USTAR initiatives have a 518 sf DC & expects to locate at (N) Temple → Data Warehouse has over 200 feeds and several terabytes of data	
2.02	UU Desires "World Class Data Center" → Consider Modular MEPFP and/or IT topologies & systems - Not fed from UU Central Plant - Minimum Tier III - No planned outages desired, NO Single Points of Failure - Concurrently maintainable, fault tolerant → Pursue LEED considerations: solar panels, heat reclaim, water efficiencies → Virtualization – Reduce (E) UU power profile; decrease over time to save $ - Almost at the point of requiring proof that systems must run on standalone equipment to not have application running on a virtual instance → www.it.utah.edu contains information on plan and focus	Info/TBD
2.03	UU Migration → Ops intends to migrate equipment from (E) sites (7) a "piece at a time" → "Remote Hands" cabinet installations (Facilities/Ops installs equipment) → Move scheduled for October 2010	Info
2.04	BOD Outline Interview (Facilities)	Info/TBD

Item	Discussion	Action/Date
	2.1 Facility Location Considerations → Wind - Large ticket item, once in a while. Need to protect gear from airborne debris (garbage, tree scraps, etc.) → Flooding – 100 yr issue; currently on lower part of valley; desert climate: not a flash hour rain issue; deal with high snow pack that melts → Explosion – N/A for storage onsite; hazardous materials off highway; no history of combustible off-gassing → Fire – (E) Temple FD access is OK; no none special UU FD requirements → Earthquake – (E) Temple at bottom of valley ~ subject to liquefaction → Hail/Snow – Thawing/refreezing ice/snow on roof may impact exposed equipment. No dramatic elevation changes in paving (wheelstops, bumps, depressions…that could be struck by snow removal equipment); snow removal is contracted; hail not a known issue → Other – Dry climate, dusty air (summer prevailing winds) → Utilities – Electrical: Overhead lines; Telco overhead/manholes; street city reads of meters → Surrounding Areas – Consideration for adjacent properties at NE corner: upgrade existing perimeter along adjacent properties? → Site Access, General Features – Site perimeter fencing with gated access: remote, card-in/out; segregated site access points (visitors: colo, visitors, delivery, refuel) and (employees: facilities, it/ops); site fencing: "candycane" shape or similar extra security level at top of fence; screened equipment: access through mandoor if equip gate is oversized; General features: cameras: perimeter of building, controllable cameras at ingress/egress points; cameras to cover general site (non improved/no equip areas); cameras in equip yards 2.2 Infrastructure Design Considerations → Essential Facility – TBD (may be driven by Colo partnering) → Redundancy – See other sections → Fault Tolerance – See other sections → Maintainability (Concurrently Maintainable) – See other sections → Facility Growth Strategy – Modular, see (ACS, OIT, HPC…etc.) 2.3 Raised Floor Design Considerations and Coordination → Minimize # ramps; consider lift options in lieu of over height assembly elevations; no applied trim, min of 4 pullers per space, grommetted / airflow blocked openings, loading/dissapitive coatings TBD per User Reqt's, 20% overstock of panels; aisles/rows designated with nomenclature (across entire facility, not just single area); crash rails at exposed equipment /conduits and pipes along aisles 2.4 Electromagnetic Interference → Not intended for satellites on roof; TBD per user requirements	

Item	Discussion	Action/Date
3.00	**Space / Capacity Requirements**	
3.01	(E) Existing UU 7 Data Centers (DC) are fully built-out (out of footprint capacity) (E) 15 sites: currently ±32,000 sf raf DC; colleges have their own sites → Seeks to consolidate down to few or 1 DC - Library recently opened a New Data Center; Komas DC may vacate → Each (E) DC currently has own change controls & scheduled maintenance	Info
3.02	Relocation/Migration → Forklift possible for relocation of equipment	Info/TBD
3.03	(E) Temple Facility: → Former, partially improved (for DC use) WorldCom/MCI site → ± 75,000 sf shell → 3 abutted buildings, "A-C", from North (oldest) to South (newest) → Building A - Medical Records Storage TI Design at approx. 95%; set to occupy → Buildings B and C intended to house UU Data Center	Info/TBD Retrieve Documents
3.04	(N) Temple Facility Interior Spaces and Requirements → Entry Lobby – Immediate Access to Security; chairs only → Pre-Security Office – Work/conf room with table & chairs → Security Office – Immediate Access to Entry Lobby; Custom millwork: cabinets, countertop, 2 levels of monitors above + independent work table; no microwave/fridge; security equipment elsewhere; closet within for minor supplies → Facilities Manager's Office – white board → Conference – "Situation Room" for Facilities, guests; white board; table, chairs; A/V in ceilings/walls, not furniture → Open Office – minimize (Facilities use) → Data Center (DC) – Separate entries for each of 3 business units (BU); Hard wall each BU; consider Vision Panel(s) from Corridor into space(s) to avoid bringing visitors into DCs → (DC) CORE – Unsecured public aisle way RAF and/or CLG panels → (DC) COLOcation – Secured public aisle way RAF and/or CLG panels; caged/secured data storage → (DC) High Density Computing (HDC) – → NOC – TBD Need ~ (E) off-site NOC with "view" of (N) Temple ok? → Meet-Me – TBD Need (COLO?) → Loading Dock (LD) (interior) – Secured Storage (hardwalled, not caged) → Facility Storage – MEP + Janitor supplies → IT Storage – Location TBD: possibly banked or direct access from each DC → Break – Refrigerator, microwave, sink (hot and cold): central table, counter (limited cabinets); open trash; vision panel into space; white board → Restrooms – Unisex accessed from Pre-secure area; bank of M/F inside secure perimeter; Unisex shower (1 min) with lockers (10)	Info/TBD

Item	Discussion	Action/Date
	→ Roof – requires permanent access: consider dual exterior building ladders (2 level roofs) in lieu of interior stair to roof (N) Temple Facility Exterior Space Requirements → Bike Storage – Covered & impact protected, not conditioned, lockable, motion-lit, lockers inside → (F) Guard Booths – Consider stubbing conduits to possible (F) guard booth locations at site entries	
3.05	(N) Temple Facilities Operations (OPS) → Up to 12-20 MEP/IT staff; combined staff for all DC operations → CHCP installs all HPC hardware, even if purchased separately by researchers → Operations Security and Compliance supports both OIT and ITS → Facility to be considered "Lights out DC", with minimum Admin spaces → UU seeks detailed MOP to be provided for long term O/M	Info/TBD
4.00	**Civil**	
4.01	(E) Temple → Verify site "hazards", especially vs (E) top of slab elevation (flood zone?) → Grade parking, paved asphalt with concrete curbs, XFMRS (2),	Info/TBD
4.02	<u>BOD Outline Interview (Facilities)</u> 4.1 Site Development and Site Work → Secure site first to allow for on-site storage (manual at min first) → Traffic Control – Gated sliders; site fencing/circulation considered of impact protection within and external to perimeter security → Site Protection and Access – → Parking and On-Site Circulation – Paving: flat non-obstructed surface; striped parking indications (not adjacent to building); designation for visitors and employees; pole mounted/paving painted; no guard house at perimeter; manned booth at loading dock area; gated loading dock similar to existing with mandoor; removal of exterior hardscape in non-used (West building) areas → Exterior Infrastructure Areas – Ganged yards ok → Public Utility Requirements, Access – TBD with Utility → Site Accessibility (Path of Travel) – To Code (ADA) → Walks – Covered at building, possible vestibule with snow melt grate → Drainage – Catch basins ok; look into OD at exterior walls draining to grade…verify grade drainage requirements at all of building perimeter; segregated fuel oil spill provisions (drains) → Landscaping – Native plants (conserve water)/groundcover; maintain irrigation systems → Soil Conditions – N/A 4.2 Environmental impact assessment → Air – Particulates	Info/TBD

Item	Discussion	Action/Date
	→ Water – Potential shortening of supply → Outdoor Noise Criteria – Enclosures at equipment to meet code; → Other – Solar energy TBD	
5.00	**Structural**	
5.01	SLC is seismically active: confirm (E) Temple completed work	Info/TBD
5.02	(E) Temple: → 18" conc reinforced (shotcrete) exterior walls → 8" conc floor slab → steel roofs (heavy load) structures (over Buildings B & C ~ can hold cars)	Info/TBD
5.03	<u>BOD Outline Interview (Facilities)</u> 5.1 General – Level of design (ie: essential in colo) 5.2 Structural Design Criteria – Code/AHJ minimum 5.3 Seismic Considerations → General – → Base Isolation – → Floor Isolation – → Other Seismic Considerations – Suspended loads: overhead inclusive of all systems overhead; housekeeping pads preferred where clearance acceptable → Seismically Rated Data Processing Cabinets – Cabinets even at RAF down to slab anchorage → Equipment Isolation Bases - Not likely Site Erection Considerations, Tolerances – TBD	Info/TBD
6.00	**Architectural**	
6.01	(E) Temple → Exposed shotcreted intr of extr walls; UU requests them furred (smooth finish)	Info/TBD
6.02	<u>BOD Outline Interview (Facilities)</u> 6.1 General → TBD: Colo tenant requirement driven (DOD) meet most stringent requirements 6.2 Base Building Requirements – Campus Design and Construction UU standards (to be provided by UU) 6.3 Roof System Requirements – Alt to stair: Ladder access in two stages, threat prohibitive from grade, interior or exterior, hoist to accompany 6.4 Data Center Requirements – No additional building standard → Space Planning – → Equipment Layout –	Info/TBD

Item	Discussion	Action/Date
	→ Access for move-in tied to relocation equipment – Common loading dock area then to access interior header collector corridor with common security review → Space Access – Minimized entrances: common header corridor at east ok; segregated entry to each of the BU's determined → Fire Rating – 1-hr min to critical + MEP, + perimeter of Admin → Partitions – Campus construction, economy, strength considerations → Doors, Frames and Hardware – Campus Standards; automatic with chain overheads… → Vapor-Tight Requirements – Air controlled data environments, applications TBD per mechanical → Acoustical Ceilings – No ceilings in DC; Admin yes; 2x2 preferred…all per campus standards → Access Floor – No (E) standard per campus standards; dictated by DC pod; integral trim only; min (4) pullers per space → Finishes – TBD in DC, all else campus standard → Specialty Millwork (NOC/ITCC/Call Centers/Trading Floors) – TBD per allocation of these spaces → EMI/RFI Shielding – None at this time 6.5 Storage/Staging – Loading dock door into non-visible pallet accessible secured storage; larger separated IT vs MEP/Facilites storage (only 1 large location for storage); add 1 storage for SEC space, can be caged; 1 at prescreen area (4 types id'd) 6.6 Office/Personnel Area Requirements – Campus standards → General – → Signage – → Walls, Elements, Flooring, Finishes and Accessories – Secured panel assemblies in interior common space, and prescreened lobby spaces (hard lid with mesh at prescreened spaces); corridor impact protection in circulation corridors and on doors to 48" (as possible) → Vision panels – Into storage, no viewing window into staging 6.7 Soundproofing – DC loudest and shall be segregated; office admin spaces INDIVIDUALLY sound separated → General – → Noise Reduction – → Noise Absorption – 6.8 Accessibility Considerations – Meet it → Parking – Wheelstops bad → Walks –	

Item	Discussion	Action/Date
	→ Ramps – → Entrances – Vestibule at main entry with recessed pit for snow off at door → Doors and Doorways – → Stairs and Stairways – → Wheelchair Lifts – Consider in lieu of long ramps at DC (HPC) → Public Telephones – Not unless required → Toilet Rooms – → Drinking Fountains – EWC Water cooler provided (hot and cold) → Identification, Signage – → Warning Signals and Hazards – → Flooring – Hard impact flooring in ALL aisles for circulation → Controls – 6.9 Ancillary Space → Staging, storage (at LD, then internal~MEP + Jan, and IT for DCs) → Situation/conference room (1 big one) 6.10 Egress Compliance – 6.11 Cages and mesh enclosures – Slider doors	
7.00	**Mechanical**	
7.01	Tier Review (Uptime institute): → Tier 2 is N+1 equipment with a single direction of chilled water feed → Tier 3 is maintainable. N+1 equipment and piping loops that allow chilled water feed from two directions in order to isolate any piece of equipment while feeding all others → Tier 4 is fault tolerant. 2N or 2(N+1) chillers with 2 separate piping loops	Info
7.02	Heating Considerations: → Hopefully a boiler won't be needed with the waste heat available from the data center → There is natural gas at the Site	Info/TBD
7.03	<u>BOD Outline Interview (Facilities)</u> 7.1 General → Systems Design – Tier III min (concurrently maintainable) ~ basic N+1 system → Base Building HVAC Design Conditions – → Ventilation and Exhaust – TBD → Support Area Design Load – TBD 7.2 Building HVAC Systems Criteria → HVAC Energy Utilization – → Cooling with Outside Air (Economizer Cycle) – Used in Winter the most, however, may not pay; Nothing against this concept	Info/TBD

Item	Discussion	Action/Date
	HVAC Systems – 7.3 Data Center HVAC Systems → RAF systems in (E) and (E) MEP familiar with RAF → Design Criteria – 72 deg now, 45-50% humidity → Cooling Load Requirement – 1.1 MW as start point; may triple for final Build-out (4-5 MW possible) → Data Center Air Conditioning – 7.4 DC Air Distribution – Hot/cold aisles: no current issue with 'hot' of hot → Supply Air – CRAC units now; currently running glycol cooling equipment, downflow units → Return Air – No use of overhead plenum returns in (E) → Supplemental Systems – If 20kW cabinets (HPC area): Overhead cooling system mounted to cabinets, with return down the back…don't provide humidification; APC liquid cabinets: reactive matrix to what's provided 7.5 Ventilation – Clean room standards 7.6 Filtration – Particulate problems in winter, review O/M 7.7 Cooling Plant – Air/Water shall be centrificals; best as modular expandable to match with final build; Exterior (outdoor, rooftop) plant → Air-Cooled Equipment – Cheaper, if design can cool load; effective to 105° → Fluid Cooled Equipment – → Chilled Water Equipment – Higher $, more sf, piping, higher cutoff top deg; best as modular 7.8 Cooling Towers/Fluid Coolers – TBD → Cooling Towers – TBD → Fluid-Coolers – If no outdoor air economizer, use fluid-coolers → Condensers – Part of systems discussed 7.9 Piping Systems – Do not let block air flow → Chilled and Condenser Water – → Refrigerant Piping – 7.10 Heat Recovery and Energy Conservation – Waste heat (from DC over) → Ventilation Energy Recovery Devices – TBD → Liquid-To-Liquid Energy Recovery Devices – → Liquid-To-Liquid Air Energy Recovery Devices – 7.11 Liquid Detection System – (E): Rackpot system underfloor around mech	

Item	Discussion	Action/Date
	units and at low points of floor 7.12 Thermal Storage – HP to propose a tank + sf 7.13 Data Center Humidification – All year required, 10-15% → Humidification at the CRAC Unit (local) – → Central or Stand-Alone Humidification – → Humidification Systems Alternatives – 7.14 UPS, Switchgear Rooms – (E) Dedicated critical cooling equipment 7.15 Battery Room – HP: cooling air + ventilation (required) in space 7.16 Mechanical Controls – UU on JC now, DCs are not…some equipment is 7.17 Plumbing – Keep water out of DC → DC Condensate Drain – → DC Under Floor Drainage – DC Floor drains not preferred; no known gas creep; Sump can be provided in DC spaces for discharge → DC Overhead Systems – → Domestic Hot and Cold Water – flash heaters at sinks → Sanitary Fixtures – Number ~ code minimum; shower required (1 min unisex); manual controls; floor drains 7.18 Standby Generator Mechanical Systems – Min 48 hr runtime; trucked in with dedicated response due to state qualifications → Diesel-Fueled Emergency Generator – → Fuel Oil System – above ground storage tank; manifolding TBD; no convaults with pumps (belly tanks keep gen above grade) → Fuel Treatment System – Heater in tank, polishing TBD → Rainwater Containment – Canopy at necessary locations → Refill, Spill Containment Procedures – All level (machines, piping, tanks, refill): containment, plan, mitigation, alarms	
8.00	**Fire Protection**	
8.01	(E) sites dry pipe + FM200	Info
8.02	<u>BOD Outline Interview (Facilities)</u> 8.1 General – Halon, preaction in use currently 8.2 Detection Systems – 8.3 Annunciation/Control, Telecommunications (Phones)	Info/TBD

Item	Discussion	Action/Date
	8.4 Air Sampling – UU to provide vendor → Detection at floor and at ceiling level; 2 station positive to go into alarm, with 30 second delay; UU seeks direction on how to "compartmentalize" to minimize impact of alarm 8.5 Suppression – Consider gas → FM-200 – Yes → NAF S-III → Inergen – Space consuming → Saphire → FE-13 8.6 Sprinkler Systems → General – → Fire Protection Materials and Equipment – 8.7 Building Fire Alarm System – No UU standard currently; on City grid for alarm response → Fire Alarm Design Criteria – → Standard Fire Sprinkler (Wet System) – Only if cost effective for admin/support spaces → Pre-Action Systems – Dual interlock in critical spaces (MEP/IT + DC); in all zones, under RAF tbd 8.8 Fire System Procedures → Fire Education and Training – Nothing unique about UU standards for fire codes; UU to provide campus standards for incorporation 8.9 Data Center Construction Standards 8.10 Automated Tape Libraries – Vaulted (floor, ceilings, walls); preaction + gas 8.11 Fire Extinguishers – On wall in secure areas, in cabinets in COLO, on wall in Pre-Secure area →Type of Extinguisher: content / area shall drive protection (type) 8.12 Storage Standards → Data Storage → (Paper) Storage Rooms – TBD per Records Storage + industry recommendations → Flammable Liquids – None (other than MEP) expected	

Item	Discussion	Action/Date
9.00	**Electrical**	
9.01	General → (E) UU sites/buildings only approx 30% separately metered → (E) ±45 yr old M/E infrastructure - 2N PDUs, N+1 UPS, N Gen, has ATS bypass	Info
9.02	Service: Rocky Mountain Power @ (E) Temple → (E) Utility: 7x outages for up to 7 hrs in past year → UU seeks "cost cap" (RMP surcharges) inquiry with Utility → (E) Temple: looped vs radial service to site? 2x incoming? 1 is overhead → (E) Temple has 2 XFMRs, only 1 is energized (Utility Owned)	Info, TBD
9.03	(N) Temple: → UU interested in $ vs 3x power level planning/tiering → Metering (levels/locations TBD) likely a key component in $$ generating	Info/TBD
9.04	UU to provide current (E) DCs' Power usage information: Estimated at 1.1MW	UU/TBD
9.05	<u>BOD Outline Interview (Facilities)</u> 9.1 General – Preferred Modular expandable topology 9.2 Primary Utility Service → Rocky Mountain Power 1.4MW to site (as discussed in CPrawdzik's email) 9.3 Secondary Service Transformer – yes 9.4 High voltage Automatic Throw-over switch – acceptable 9.5 Medium Voltage Distribution – as required 9.6 Main Service Switchgear – as required 9.7 Standby Power System → Types of Generators – Diesel → Generator Sizing – sized to load → Generator configuration – N+1 → Generator performance characteristics – TBD 9.8 Transfer of Source – acceptable → Motorized Breaker Transfer – acceptable → Automatic Transfer Switch (ATS) – with isolation Bypass (for concurrent maintenance) → Generator Peaking or Load Shedding – not required → Standby Load Prioritization – not required, large loads not anticipated → Transfer system Capabilities – as noted for ATS 9.9 Test Switchboard – as required	

Item	Discussion	Action/Date
	9.10 Electrical Distribution – as required 9.11 Transient Voltage Surge Suppressor (TVSS or SPD) – yes 9.12 Uninterruptible Power Supply – yes → Criteria for Analysis – → UPS configuration – N+1 → UPS characteristics – 9.13 Battery Systems → Vented (Wet) – longer life/higher initial cost – yes → Valve Regulated (Dry) – TBD → Rotary Hybrid – no 9.14 Battery Monitoring – yes 9.15 Battery Spill Containment – as required by Code 9.16 UPS Bypass Distribution – yes 9.17 UPS Distribution Panelboards – yes 9.18 Line Voltage Transformers – yes 9.19 Static Transfer Switch – no 9.20 Data Center Power Distribution Units (PDU) – yes in Data Space 9.21 Computer Power Center – yes/RPP in Data Space 9.22 Remote Distribution Cabinets (RDCs) – n/a 9.23 Line Voltage Distribution – n/a 9.24 Grounding – yes, Code and DC zero reference ground 9.25 Lightning protection – yes 9.26 EPO – yes if required by Code – no otherwise 9.27 Lighting – → (E) Temple: leave existing HID outdoor as is; consider leaving (E) indoor	

Item	Discussion	Action/Date
	metal halide to remain (even if above suspended ceiling) → (N) Temple: fluorescent strip fixtures (DC), balance TBD 9.28 Harmonic Distortion – K rated transformers 9.29 Convenience Receptacles – yes at interior walls of DC 9.30 Electrical System Monitoring – yes BMS 9.31 Underfloor Cable Management – Cable tray - overhead (flooding) 9.32 Fuel Cell Technology – no 9.33 Seismic Bracing – yes 9.34 Temporary Generator Connection – yes 9.35 Load Bank → Yes, if within budget → Connections: for Generators – yes; for UPS – yes	
10.00	**Security**	
10.01	(N) Temple: 3 pod security: Hospital/Core, Colo, HPC → Colo will force hardened partition/cage separation from Core & HPC → anticipate common Facility entry with controlled passage to each of 3 areas (E) UU sites + (N) Temple: access will be badged	Info/TBD
10.02	<u>BOD Outline Interview (Facilities)</u> 10.1 General 10.2 Security Design Objectives – No blind spots outside of extr. Fencing, but fixed ok; cameras at all site access points; cameras at yard access points, not necessarily at secured (fenced equipment); at common/visitor/vendor/colo ~ all locations monitored → Functional Objectives – → Site Access – Vehicular: badged with call button; man/bike: badged (keyed Facilities) with call button → Utility Yards – badged (keyed Facilities) with call button → Building Entrances – Mantraps; physical barrier with slider, human to human view, with small lobby; badged (keyed Facilities) + biometric-like with call button back to local station → Exterior Signage/Accessibility – Minimize, conceal UU identity → Perimeter Lighting – Per code, per maintenance, per Security monitoring provision	Info/TBD

Item	Discussion	Action/Date
	→ Fencing – required: at site perimeter, interior of site at yards ("candycane") → Other Access Points – TBD Utility; walk-up/biking → Concealment Areas – Limit general/eliminate at critical → Critical Communication Facilities – MPOE (IT sensitive): badged (keyed Facilities) + biometric-like with call button (local) → MEP: badged (keyed Facilities) + with call button → DC: badged (keyed Facilities) + biometric-like with call button 10.3 Space Classifications → Public Space – → Camera with sound – Yes, TBD locations, see Space descriptions → Interior Space – Cameras: Building ingress/egress locations, corridors, common spaces; ingress/egress of interior critical spaces (MEP + DC); cameras not necessarily required at interior equipment of MEP spaces; DC COLO: no concealed common space; DC CORE/HPC: circ aisles → Restricted Space (High Security Area) – Tape Storage Room: badged (keyed Facilities) + biometric-like with call button, sound detection; liquid tight (provide leak detection) horizontal (clg, flr) TBD on wall/wall ht; 2-hr voluntary fire resistive enclosure min 10.4 Space Security Criteria → Lobby – "Pre-Security" area; Mesh/masonry at interior of lobby → Loading Docks – Cameras: hand/zoom/tilt all LD bays, intr/extr, see other gates for access requirements; run conduits for future booth connections at site access/loading dock points → Data Center – Mesh/masonry at interior of DC COLO demising partitions and at DC that borders corridor → Security Monitoring Station – badged (keyed) → Break room – badged (keyed) 10.5 Architectural Features → Walls and Partitions – Mesh/masonry to separate secure/pre-security → Windows – Exterior: none existing; Interior: 1" bank security; Viewing panels at DC COLO only (bullet resistant in lieu of perimeter metal detector scan) → Doors – Entry door: not glazed, not extra heavy; extra heavy door at loading dock into building; MOP to control moving delivered goods into facility and/or extra heavy door to limit delivery visitors to access storage of received goods 10.6 Controlled Access System (CAS) → General → Card Access – Badge in/out of all DC, Tape, MPOE spaces; have Johnson control now, look at others	

Item	Discussion	Action/Date
	→ Biometric Authentication – 10.7 Alarm Systems – UU standards; local and relayed offsite 10.8 Digital Closed Circuit Television (DCCTV) – Same as Alarm systems 10.9 Equipment Cabinet Security – Key access in DC COLO to cages and cabinets 10.10 Security Procedures – TBD with others; written procedures, posted locations TBD (local and remote)	
11.00	**Technology**	
11.01	UU: pursuing Fed funds for local (UU) or municipal (SLC) dark fiber ring (MAN) → Connectivity TBD or to be provided at Temple → (E) Fiber right-of-way along 8th Street → Qwest Communications indicates it will NOT block fiber ring	Info/TBD
11.02	Wireless: (E) + (N) Temple + (E) UU DCs ~ Cisco (provider)	Info
11.03	Increase Virtualization as Hardware amount to decrease = increased savings → Goal: lose redundancy in services (Core-Hospital + Main UU Campus) "Virtual Desktops" ~ currently deployed through Citrix → UU seems to be accepting of virtual solutions and not requiring physical access to systems. Remote access is OK → Where to locate Hardware? Speak with Student Affairs & Dan Bowden for Hospital – Possible high MEP consumption	Info/TBD
11.04	ACS, OIT, ICS, Media Solutions, Libraries: currently spending $ to upgrade → Moved to a centralized Mail Service	Info
11.05	Information Tech Council (ITC): → Represents major groups: IT, UU, EDU → Handles (Individual College) Community Services: Network, Antennas, Email	Info (UU to provide Org Chart)
11.06	Single Points of Failure current in UU IT operations	Info
11.07	<u>BOD Outline Interview (Facilities)</u> 11.1 General 11.2 Codes and Standards 11.3 << *insert* >> City and << *insert* >> County Background 11.4 Design – Consider long term for fiber network 11.5 Utility (Carrier) Service – (E) Temple has carriers: Qwest, others TBD 11.6 Cabling System Infrastructure	Info/TBD

Item	Discussion	Action/Date
	→ Building Entrance Facility Room (BEFR) → Main Communications Room (MCR) → Computer Room (CR) → Telecommunications Room (TR) → Backbone Cable System → Horizontal Cable System → Station Outlet → Cross Connection and Interconnection → Patch Panels → Insulation Displacement Connectors → Building Entrance Protection Device (BEPD) → Grounding 11.7 Design and Installation Considerations → General → Cable System – Overhead data cabling in fiber (all) → Communications Raceways → Seismic Bracing → Plenum Installation → Only electrical → Labeling and Marking – Yes! → Final Testing → Cable Separation From Power Wiring – In different planes in DC → Meet-me – Likely required (COLO) 11.8 Twisted Pair Cable Recommendations – Cat 6 now 11.9 Multimode Fiber Optic Cable Recommendation – UU standard 11.10 Fiber to the Desktop – Not necessary 11.11 Wireless Systems – Yes, Cisco 11.12 Equipment Security and Monitoring	
12.00	**BMS – Building Management System**	
12.01	UU: Automation is critical; consider monitoring of the following: → Plug strips at server level / Consumption at Cabinet/Rack level → Batteries (power levels) → Mechanical Equipment	Info/TBD
12.02	UU currently employs Johnson Controls (MEP)	Info/TBD
13.00	**Media**	

Item	Discussion	Action/Date
13.01	(E) UU Media requirements not available for review at this time	Info/TBD
13.02	<u>BOD Outline Interview (Facilities)</u> 13.1 General 13.2 Design 13.3 Video Systems → Cameras – Set Conf room for Video conf (all clg mounted, not in furn) → Projection Systems – Yes in Conf room (instead of screens) → Screens, Plasma, LED, Laser, HDTV – None if projection screen duplicates → Monitors – IT preference → Lighting – → Recording (Taping) Equipment – Conf room speakers/microphones; Conf room (consider STC of 45 or higher) 13.4 Sound/Acoustic Systems → Microphones – IT preferences → Speakers – IT preferences →PA Systems – Throughout Facility 13.5 IT, Data, Telephone – Standard UU → Cabling – → Jacks – → Hardware – → Software – 13.6 Operating Centers and Control Rooms 13.7 Monitoring, Controlling 13.8 Furnishings → Hardware – → Consoles, Casework, Workstations, Seating – Custom millwork to UU standard; rolling chairs → Tables – Free, not anchored → Wall-mounted/-integrated Boards – White boards in Sec and MEP offices 13.9 Architectural Considerations → Egress – → Security – → View "Site" Angles, Elevations: TBD → Finishes (Flooring, Partitions/Walls, Ceilings) – → Partial-, Full-Height Glazing – → Doors – → Storage, Staging –	Info/TBD

Item	Discussion	Action/Date
	13.10 Power – 13.11 Lighting – UU standard in office/sec/lobby; task lighting at control millwork (under shelves); Conf Room Dimmable, 3 cones at board 13.12 Cooling 13.13 Fire Protection 13.14 Fire Alarm, Controls, Relay 13.15 Training 13.16 Maintenance, Operations	
14.00	**LEEDS / Green Design**	
14.01	Pursue Water + Emissions for energy savings – DOT (economizer, reheat, etc.) → contact Kent Udell	Info
14.02	(E) SLC: airborne particulates, water rights/desert drought climate + population to double in next 40 years	Info/Tidd
14.03	Discussed green possibilities: → Variable frequency motor drives. → Outdoor air economizer → Evaporative cooling – water is scarce in SLC → Waste heat recovery → Chilled water storage for maintainable cooling & cooling @ peak power times → Solar – will check practicality for data centers → Check ground source heat & cooling	Info/TBD /Tidd
15.00	**Testing / Commissioning**	
15.01	UU desires SOP that includes User/Facilities testing PM schedule	Info/TBD
16.00	**Schedule**	
16.01	Gnatt schedule to be provided	Info
16.01	Key Milestones/Meetings → Site Visit accomplished May 13th	Info/TBD
17.00	**Budget**	
17.01	Confirmed Direct Construction Cost budget is ±7 million dollars (Phase I?)	Info/TBD

Item	Discussion	Action/Date
18.00	**Studies**	
18.01	UU Risk Assessment ongoing	Info

Next meeting: **To Be Decided**
Location: **Conference Room**

EYP MCF will rely upon these minutes of meeting as a record of the items discussed and conclusions reached. Please advise the undersigned of any additions or corrections.

Submitted by,

Charles Prawdzik Jr., AIA
Principal Project Manager

Windows Technology Discovery Session - 5/13/09			
Name	**Organization**	**Phone**	**Email**
Earl Lewis	OIT	801.581.3635	earl.lewis@utah.edu
Steve Adams	ACS	801.581.3408	sadams@acs.utah.edu
Bryan Peterson	ITS	801.587.6095	bryan.peterson@hsc.utah.edu
Mike Basinger	Marriot Library	801.581.3753	mike.basinger@utah.edu
Caprice Post	OIT/ITS	810.585.5404	caprice.post@utah.edu
Andrew Reich	OIT	801.587.0902	andrew.reich@utah.edu
Steve Carter	HP/EYP - Tech. Cons. - PM	312.343.9535	scarter@hp.com
Rob Myers	HP/EYP - Tech. Cons	312.909.1567	rmyers@hp.com

Italics indicates new discussion or corrections/updates/results/response from a prior meeting

OIT
Citrix/Exchange/Messaging

1. Growth – major upgrade 2003/2007. 30% growth in past 12 months.
2. Small virtualized environment. 25% growth this year. 25 machines.
3. Active/Active clusters. Both clusters in the same data center. RTO is 3 days to a week – customers understand this is the case.
4. Inmagic – software solution extended synch and backups. Deltas will be stored in Richfield.
5. 88,000 email boxes. Expect number of email accounts to stay relatively static, but do expect storage needs to grow.
6. (5) – 1855 Dell Blade Chassis – 10 blades per chasis. One common vendor initiative.
7. Dual source power circuits designed and required.
8. Voice mail – 1 current standalone Intel boxes. 3 in the future. Future blade servers.
9. No standard for power per rack.
10. Avaya systems not likely to move.
11. 6 Active directory servers – not expected to move.
12. Inventory control – bar coded medical system. Exchange systems not tracked - manually tracked.
13. Growth in utilization of existing email boxes.
14. Upgrades in Exchange have dictated different hardware requirements.
15. 3 year refresh rate. – Next refresh planned to be placed at New Data Center. Existing equipment would move to Richfield for DR.
16. 3 copper / 2 fiber connections per enclosure.
17. Development / Test – Retired production becomes development. Test environment – VM. Development environment needs UPS and generator – Test is really pre-prod so should run in new data center.
18. Do not need to touch equipment if competent data center managers/staff.

19. Physical access to machines is required test/dev for patching and power failure type of events. Remote power on of hardware required.

ITS

1. RTO – not well defined
2. RP0 – not tolerant to data loss. Data loss is major issue with Health Sciences. Longer down time is acceptable as opposed to losing data
3. Inventory tracking – HP Insight Manager. Altiris asset manager.
4. Blade chassis – 2 X 10Gig. ILO for remote management. Team NICS standard for standalone.
5. Refresh rate – 3 to 5 year refresh rate. 5-10% annual growth.
6. ESX – 200 guests. Looking at HyperV. Will have a mixed environment.
7. 10 – 1 VM. BL25PG1. BL490CG6. ESX growth.
8. Virtual Citrix environment – using Zen. BL465G5.
9. Rate of virtualization – rapid at first. Slowing down.
10. SSL encryption and Java based apps do not do well VM.
11. Database performance issues in virtualized environments.
12. 2 – 4 enclosures per year.
13. Looking at Dynamic power capping.
14. Rack standard – APC 30 inch X 48 inch depth.
15. Densities – P class – 5 enclosures.
16. 380 – 20 servers per rack.
17. 360 – 40 servers per rack
18. C – 4 enclosures per cabinet.
19. LINUX machines – 4 network connections per server.
20. Copper – Gigabit over copper. Copper 10 GB – 10 meters.
21. 5 – 6 times per year. Power loss to major systems.
22. Power issues are a significant impact to the productions IT systems many times a year.
23. Storage data requirements.
24. Network failover, bandwidth and redundancy critical path issue.
25. VMotion – DRS running in automation.

ACS

1. Application data on the LINUX/Oracle systems.
2. Strong move to virtualization.
3. 60 guest OS instances.
4. SUN X4600 8 Way – 48 GB RAM.
5. Chronos – moving to VM.
6. Average P-V ratio: 15:1. 20:1 future planning.
7. VMotion – utilized.
8. 5 network drops + Fibre Channel connections.

9. Dell racks for Dell equipment – Sun equipment in Sun racks.
10. KVM remote access.
11. Remote access that works is fine.
12. No rack space standards – grouped by function.
13. Campus SharePoint services initiative, estimate 3 to 4 chassis for new SharePoint services support.
14. 1 SUN chassis per year growth. (20 guest instace growth per year)

Marriot Library

1. Just built a new 2700 SF datacenter.
2. May not need space in downtown datacenter.
3. SAN – mirrored. Would like to locate one cluster in other data center.
4. Routing of info to library has some unique out of bandwidth requirements.
5. Dell 2950 boxes.
6. SAN CX400 – 17 enclosures (750 or 1000 GB drives). Connected to (2) 6650 Dell server. 14 network connections off of SAN
7. Virtual – ESX running on the 6650 servers.
8. V Motion not utilized, but tested.
9. Network servers – 3 netwok connections per server.
10. New data center – with current planning should meet 3-5 year growth.
11. All Dell systems – rack systems.
12. Mac group. Cluster systems for certain users.
13. Redundant PDU, UPS and N+1 CRAH units. Cooling issues have been an issue in this new DC. Lost cooling in the past due to misconfiguration of the Leibert systems.
14. Grants could drive unexpected growth in data center.

Next meeting: **N/A**
Location: **N/A**

EYP MCF will rely upon these minutes of meeting as a record of the items discussed and conclusions reached. Please advise the undersigned of any additions or corrections.

Submitted by,

Robert Myers
Senior Technology Consultant

Unix Technology Discovery Session - 5/13/09			
Name	**Organization**	**Phone**	**Email**
Earl Lewis	OIT	801.581.3635	earl.lewis@utah.edu
Mike Basinger	Marriot Library	801.581.3753	mike.basinger@utah.edu
Andrew Reich	OIT	801.587.0902	andrew.reich@utah.edu
Chad Lake	OIT/ITS	801.230.3086	chad.lake@utah.edu
Corey Pedersen	ACS	801.581.3637	cpedersen@asc.utah.edu
Tim Richardson	ACD	801.585.1152	trichardson@acs.utah.edu
Bryan Peterson	UEN	801.585.7789	bryan@uen.org
Bryan Peterson	ITS	801.587.6095	bryan.peterson@hsc.utah.edu
Steve Carter	HP/EYP - Tech. Cons. - PM	312.343.9535	scarter@hp.com
Rob Myers	HP/EYP - Tech. Cons	312.909.1567	rmyers@hp.com

Italics indicates new discussion or corrections/updates/results/response from a prior meeting

OIT/ITS

1. SUN Equipment - T2000's 5 years at end of life.
2. SUN racks.
3. Virtual infrastructure – Rack mount systems. Blade environment running Linux.
4. DB (5) – 490's. 5 years EOL. Not virtualized
5. Virtual environments – growth available for environment. SUN 4600 standard environment. 64 GB of RAM.
6. 500 Web Servers.
7. Vendor support always an issue for virtualization.
8. 30 SUN 4600's a minimum base virtualization hardware environment for future. 50:1 ratio – small load allows for large ratio.
9. T2000s are currently hosting 250 websites.
10. Assume 25-30 virtualized machines day one – assumming 50:1 ratio achieved.
11. Network 10GB x 2 maybe X 5. 4 1 GB connections.
12. Webserver's, identity management, time, DNS, DHCP. VM to campus. MYSQL, Oracle hosting. Web applications hosting. Content management – Vignette. WIKIS, list servers, blogs, Google APPS, search and web trends analytics – all functions would move to new data center.

Utah Education Network (State level support)

1. Provide state-wide services K-20.
2. SUN mostly. Blade 6000 chassis. SUN E 6000 servers. Virtualized into zones. 16 : 1 virtuization.
3. Blackboard Vista primary application environment.
4. Oracle running enterprise class hardware – moved to new DC. SUN E4900.
5. T 2000's still on floor.

6. (2) Current T 5240 new SUN platform. 2 per year.
7. (4) Current E 6000 servers. 1 per year starting 2010.
8. Do not plan on vacating EBC-UEN data center only. Larger data center is the UEN. Half the gear would move to virtualized platform.
9. Primarily host Web apps, blackboard vista (course management), video management, education management.
10. Outage windows – 7 X 24 type of environment. Due to statewide user base.
11. Management – power strips serial connections via IP to serial consoles. Advocent remote management capabilities. Lights out management.
12. 4 copper GB per blade. 40 GB per chassis. Fiber channel SAN connectivity.
13. SAN would be replicated from UEN DC to new data center.
14. Fiber channel and Ethernet switches per rack.

ITS

1. IBM servers. P series approx. 30 machines.
2. House SUN 250, 420R, 280R, 210,240, 440, 880, T 2000, 4000, 5000.
3. 40 LINUX – non RISC based platforms.
4. IBM consoles. New HMC's. Raritan remote access consoles.
5. Lights out management.
6. Refresh rate 7 – 10 years.
7. Virtualization – minimal virtualization now or planned. No virtualization on SUN platforms.
8. 10 % growth records.
9. EMR – next from Cerner Millennium to Epic Care – IBM/Cerner to Epic Care/SUN. Just purchased (2) M 5000's.
10. Network connections – LINUX bonded
11. Oracle rack – 10 GB environment.
12. Storage availability major requirement
13. Running Cerner, Revenue cycle, Oracle, SYBASE, Cache, Data warehouse.
14. Running Veritos tape backup systme
15. Running PACS tape archival systems.
16. Budget planning and purchasing software hosted.

ACS

1. Primary applicatins: Portal, financial and HR/payroll/Student Administration.
2. Other applications: Degree audit, chronos time and attendance.
3. Directory services, proxy and service management. NFS services.
4. Java application – custom coded environment running against above applications.
5. Current footprint – 30 racks. 22 servers and 8 for storage.
6. 3- SUN 6900, 2 – 4900, 6-8904, 8-5240, 20-5220, 1-T2000, 5-8904, 15-X4200
7. 6900/4900 Databases.

8. 8904 – Development systems, netbackup and LTO drivers.
9. 5240 – (2) dev / (2) test/ (4) production glass fish (financial, HR, student and auxiliary).
10. X4200 – planview, database storage for front end Windows systems.
11. 4600's running ESX. Infrastructure Fiber channel. Current (15) 4200 moving to (2) 4600. 10:1 ratio PtoV.
12. 4540, 7210 (NFS) storage servers.
13. Data center (Park Building) may not be moving – 600 KVA CAT generator. 300 KVA UPS. Core network node for University.
14. 6 racks for collocation services for other departments. Park building document storage.
15. Optimization efforts keeping a neutral growth for space, power and cooling.
16. CRM new capability for the University. May be a growth driver.
17. Zones for development environment.
18. Crystal services / Citrix systems Windows platforms – user interface and reporting interface. If UNIX systems remain in current DC, these Windows systems would likely need to remain with backend UNIX systems.
19. 4 more rack routing/patch systems. Gig/10 gig. 2 – 10 Gig links to campus core.
20. ACS DC may be a candidate for remaining in place.
21. 10:1 ratio for virtualizaion, SPARC based virtualizaion into zones
22. 24x7 operation
23. Printing occurs on raised floor of data center, but located in separate area (not isolated by walls)
24. Interested in locating 4 racks and tape system from Richfield (current DR) to new data center.

Next meeting: **N/A**
Location: **N/A**

EYP MCF will rely upon these minutes of meeting as a record of the items discussed and conclusions reached. Please advise the undersigned of any additions or corrections.

Submitted by,

Robert Myers
Senior Technology Consultant

Project: **Data Center Improvements - Programming**
Client: University of Utah - Project 20109
Location: University of Utah Campus
Purpose: Storage Technology Discovery Session
Date: May 13, 2009

Storage Technology Discovery Session - 5/13/09			
Name	**Organization**	**Phone**	**Email**
Earl Lewis	OIT	801.581.3635	earl.lewis@utah.edu
Chad Lake	OIT/ITS	801.230.3086	chad.lake@utah.edu
Tim Richardson	ACS	801.585.1152	trichardson@acs.utah.edu
Corey Pedersen	ACS	801.581.3637	cpedersen@acs.utah.edu
Jeff Hadden	ITS	801.587.6041	jhadden@hsc.utah.edu
Bryce Dawson	ITS	801.587.6206	bryce.dawson@hsc.utah.edu
Bryan Peterson	UEN	801.585.7789	bryan@uen.org
Steve Carter	HP/EYP - Tech. Cons. - PM	312.343.9535	scarter@hp.com
Rob Myers	HP/EYP - Tech. Cons	312.909.1567	rmyers@hp.com

Italics indicates new discussion or corrections/updates/results/response from a prior meeting

ITS

1. XP 24000 (2) – Almost fully populated. (Hitachi – main storage arrays).
2. USPV – (2) Healthcare, Exchange. All storage systems in one data center – looking at Richfield for DR – currently sending backup tapes to Richfield
3. Health Care Services – And OITS
4. 2 NetAPP Filers Using Hitachi Storage. USPV – AMS
5. TIER 1-3
    a. Tier 1 – USPV, XP 24000. Transactional DB, EMAIL, critical informaton
    b. Tier 2 – Modular EVA, AMS 1000, 9585. Development/test,
    c. Tier 3 – SATA direct attached to USPV Backups
6. Tiered Storage Manager – Storage manager application from Hitachi.
7. Storage Tek 8500 tape libraries
8. Fiber Storage – 6140 Fabric (A and B fabric) – 32 edge switches. Core to edge design.
9. Brocade DCX upgrade in August.
10. 20% DDUP recovery.
11. New Projects – 0.5 Petabyte current. Doubling of data every 12 months.
12. 12 months or less – will be out of current growth capacity of existing storage frames.
13. 5 -7 years life span. Outages for moving data are very difficult in the health services area.
14. Research growth is accounted for in the growth numbers. But new TIER 4 data storage will be implemented for research only data storage. Will be discussed in the HPC stakeholder meeting.
15. In the Komas Building – some equipment is in Richfield.
16. ISCSI – TIER 4. Tier 4 will not be backed up or redundant. Main use is computational data that can be recreated.
17. SharePoint Services, Unified messaging, Video storage will add more storage requirments.
18. Genomics, Radiological growth higher growth. Data storage may have to be close scanning machine. May not be able to move large data over distances to a remote datacenter.
19. Small datacenter inside hospital
20. Talking about separate SAN fabric just for genome issues

21. Use Citrix for doctors viewing images
22. Moving to virtual desktops would be new requirements, not included in current growth/plan

ACS

1. Hitachi 9990
2. Can externalize another 32 TB of data.
3. 7210, 612, 6140 open storage
4. 4540 Thumper
    a. Primary use of thumper is virt tape repository at DR site 2 of 3 at richfield
5. Current initiatives – ZFS for shapshot and virtualization
6. Qlogic – 96 port switches – Storage systems interconnected through Qlogic.
7. Splitting student administration and HR systems. Will require growth in higher end / high speed data storage systems. Hitachi 9990 systems.
8. High end production systems on the 9990 platforms, must plan for growth.
9. 2 year old system.
10. External data move to other systems.
11. Utilized 70% capacity of Hitachi
12. Capacity realized in next 12 months.
13. ZFS – preferred method of new growth capacity for ACS. File system based administration compared to SAN administrative level of management.
14. 3 -5 years hardware life.
15. Dataguard and VTS copies for remote site critical data backup. May use new datacenter as DR site.

UEN

1. Full mirrored storage SAN. Virtualization.
2. Storetek 6140. Sun 3510.
3. 30 TB mirrored. 60 TB total.
4. Tiered storage based performance requirements.
5. Hitachi ASN 2100's. pair here and 1 in richfield with a storagetek 6140
6. Use dataguard for oraclr remote standby
7. Hope is to move one of the mirror nodes from UEN DC to NDC
8. Richfield – equipment. SAN fabric is all qlogic – stackable 5602 4gig fiberchannel switches
9. Richfield has older 2 gig model of 5602
10. A/B fabric – Qlogic SAN fabric. Stackable 6402 Fiber Channel Switches.
11. Growth rate – double every year. Conservative.
12. Yearly adding physical storage hardware for growth.
13. Utilizing thin provisioning of storage – 80% utilization but oversubscribed 10:1.
14. Block replication to Richfield.
15. Dataguard is utilized.
16. Datadomain – 660 VTL primary. For snapshot data mirroring. One at data center and one at Richfield

OIT

1. Network Appliance – DUAL headed 3020 –
2. 3140 to replace 3020, 25 TB usable space.
3. Providing Webservices, Media streaming 5 TB, Virtual Machine images
4. Natural History Museum – new capabilities for virtual tours, end user. Storage of data may need to be local for this application.
5. Long term goal to implement additional NetApp – NetApp MetroClustering via Fiber Channel capabilities.
6. Remote replication of NAS device. IP and Fiber Channel.

General Discussion

1. OITS and Health Services – intent to consolidate tape library.

Next meeting: **N/A**
Location: **N/A**

EYP MCF will rely upon these minutes of meeting as a record of the items discussed and conclusions reached. Please advise the undersigned of any additions or corrections.

Submitted by,

Robert Myers
Senior Technology Consultant

Network Technology Discovery Session - 5/13/09			
Name	**Organization**	**Phone**	**Email**
Tim Urban	OIT	801.556.7934	tim.urban@utah.edu
Joe Breen	Center for HPC	801.550.9172	joe.breen@chpc.utah.edu
Steve Corbato	CI Strategy	801.585.9464	steve.corbato@utah.edu
Gary Vandertoulen	ITS	801.587.2137	gary.vandertoulen@hsc.utah.edu
Jeff Hadden	ITS	801.587.6041	jhadden@hsc.utah.edu
Bryce Dawson	ITS	801.587.6206	Bryce.dawson@hsc.utah.edu
Bryan Peterson	UEN	801.585.7789	bryan@uen.org
Kelly Genessy	UEN	801.209.7459	kelly@uen.org
Steve Carter	HP/EYP - Tech. Cons. - PM	312.343.9535	scarter@hp.com
Rob Myers	HP/EYP - Tech. Cons	312.909.1567	rmyers@hp.com

Italics indicates new discussion or corrections/updates/results/response from a prior meeting

Optical Metropolitan Area Network

1. Year ago planned for optical network when considering purchase; planning is currently underway.
2. MAN – Salt Lake City optical network.
3. UUtah – peer institutions – Univ would own the fiber and manage.
4. Ring technology – (10) 10 GB Waves. DWDM up to 80 wave lengths.
5. Fiber Channel support via the DWDM.
6. May need to be able to support Infiniband. Fiber channel more possible than infiniband
7. Ethernet capability. 1.5 $MM for 4 – 8 sites.
8. Level 3 node (I2), Utah board of Regents, Salt Palace, Quest and AT&T key sites to land.
9. Quest and AT&T key sites to land.
10. 4 Wavelengths from Komas to new data center.
11. UEN – provides networking for education systems for the state. Network would provide access from Salt Lake MAN to Richfield.
12. Contemplating extending the ring on a statewide level – going down I15 towards Las Vegas and LA with extension of SLC MAN ring.
13. Commodity and research connectivity moves through UNE to public access.
14. Metro optical network to get to the new data center and major carrier facilities as well as other potential partners.
15. Thinking is University would own fiber, manage and provision
16. In negotiations to acquire fiber – timing may be advantageous
17. Ring 10gigbit per second waves to data center. Looking for upgrade 80 wavelengths dwdm. Need to run fiber channel
18. Looking at UEN network for getting to Richfield

New Data Center

1. Two entrance points of ring into facility – preferrably opposite ends of building
2. End of row or top of rack switches. Nexis series switches. Come to two aggregation switches - Nexis 7000
3. End of row – (2) 5000 per row.
4. Nexis 2K top of every rack for higher density server concentration. 10 GB length limitations.
5. Hospital today is using to HP virtual links for blade chassis. 10 GB connections back to Nexis 7000.
6. Edge Access layer switches for higher density (blade chassis) systems.
7. Concerns with not centralized fiber patch panels due to crossing IP vs. fiber channel connections, or crossing A and B fabrics at the patch panel level.
8. Hitachi and other systems – cabling from below – will need to accommodate this need. May need to isolate these systems so they do not affect neighboring cabinets.
9. Design Concept - No copper cabling leaves the row. Fiber only from row to row or back to consolidation point.
10. Copper within the row utilized for lower cost of data cabling within the row.
11. OIT and Hospitals on campus getting 4 wavelengths connecting to existing campus core. Install 2 highend routers tieing current dc to new dc. Installing 2 large firewalls.
12. Looking at end of row design or top of rack. Looking at Nexus series. End of row switches to aggregation switches for campus switches and hospital DC.
13. Looking at middle row switching for length of fiber
14. Need to consider isolation of SAN and core network
15. Seismic considerationsneed to be included in BOD, foor floor, cabinets, trays, cable, etc.
16. Some researches may have separate wavelengths to prevent intermixing
17. Need a layer 2 exchange point – will need a "meet me" room.
18. Staging area, receiving area, storage space, possibly separate between hospital and general university.

Next meeting: **N/A**
Location: **N/A**

EYP MCF will rely upon these minutes of meeting as a record of the items discussed and conclusions reached. Please advise the undersigned of any additions or corrections.

Submitted by,

Robert Myers
Senior Technology Consultant

Operations, Help Desk, User Support Discovery Session - 5/13/09			
Phillip	**Organization**	**Phone**	**Email**
Earl Lewis	OIT	801.581.3635	earl.lewis@utah.edu
Dave Huth	OIT	801.585.9467	dave.huth@utah.edu
Chad Lake	OIT/ITS	801.230.3086	chad.lake@utah.edu
Tim Richardson	ACS	801.585.1152	trichardson@acs.utah.edu
Caprice Post	OIT/ITS	801.585.5404	caprice.post@utah.edu
Kelly Genessy	UEN	801.209.7459	kelly@uen.org
Michael Timothy	OIT	801.587.0108	michael.timothy@netcom.utah.edu
Bryan Peterson	ITS	801.587.6095	bryan.peterson@hsc.utah.edu
Andrew Reich	OIT	801.587.0902	andrew.reich@utah.edu
Phillip Kimball	OIT/ITS	801.587.6262	phillip.kimball@hsc.utah.edu
Steve Carter	HP/EYP - Tech. Cons. - PM	312.343.9535	scarter@hp.com
Rob Myers	HP/EYP - Tech. Cons	312.909.1567	rmyers@hp.com

Italics indicates new discussion or corrections/updates/results/response from a prior meeting

1. Discussion has not been had on who would is filling the role of incident response management – data center operations or Helpdesk
2. Currently merging both OIT and ITS helpdesks.
3. Help desk manager would need access to data center – remote access may be acceptable depending on operational procedures.
4. Biggest concern is incident management; group would be involved in restoring services.
5. May be the need to have office space in operations center of new data center – not raised floor access.
6. Initial thoughts are at least 2 spaces for helpdesk…follow up with Glen and Brent about this need
7. UEN does not need an on site presences
8. Would like a POTS line in dc for monitoring.
9. Need a staging area for equipment, and builds, with appropriate power, network connectivity, racks, phone
10. Will need an asset manager at the new DC.
11. Medical systems, ACS and OIT can share staging space, but separate designated areas within the space.
12. 256 sun servers an example of a large shipments which may need to be housed for staging.
13. Maybe recommend stage and build (build perhaps on UPS)
14. Build area also becomes decommissioning area.
15. 1000 sq. ft. staging area
16. Build area needs desk, phones, etc
17. Need secure area to store spare parts
18. Over capacity of 100 sq.ft. right now. Using data center as a storage now.
19. Tapes go to vault – management of tape off site storage-for receiving.
20. Staging, build, storage, receiving, tape handling – needs

21. Make sure LO (e.g. Hitachi) where it phones home and needs controls in place
22. Hospital equipment is not owned by UofU. Need a third party zone if 3rd party needs to work on machines.
23. Cell repeaters throughout the space to ensure cell coverage.
24. Crash carts available.
25. RFID - because now the system for annual inventory increases risk. Stickers are on the inventory and in some cases physical equipment muist be handled to locate and scan
26. Separate staging area will need to be available for Colo for stage, build, storage
27. Place to store definitive software libraries, cabinet with locks would suffice.

Next meeting: **N/A**
Location: **N/A**

EYP MCF will rely upon these minutes of meeting as a record of the items discussed and conclusions reached. Please advise the undersigned of any additions or corrections.

Submitted by,

Robert Myers
Senior Technology Consultant

ACS Enterprise Applications Discovery Session - 5/14/09			
Name	**Organization**	**Phone**	**Email**
Bryan Harman	ACS	801.585.3516	bharman@acs.utah.edu
Joe Taylor	ACS	801.581.3325	jtaylor@acs.utah.edu
Dave Huth	OIT	801.585.9467	dave.huth@utah.edu
Mike Robinson	ACS	801.585.9077	mrobinson@acs.utah.edu
Tim Richardson	ACS	801.585.1152	trichardson@acs.utah.edu
Corey Pedersen	ACS	801.581.3637	cpedersen@acs.utah.edu
Steve Carter	HP/EYP - Tech. Cons. - PM	312.343.9535	scarter@hp.com
Rob Myers	HP/EYP - Tech. Cons	312.909.1567	rmyers@hp.com

Italics indicates new discussion or corrections/updates/results/response from a prior meeting

People Soft –Software

- Maintain
- Customize people soft.
- Off the shelf plus customized modules
- Bridges to financial systems
- Major application development – Solaris platform using Java netbeans and Oracle net designer tools. With proliferation of SharePoint and SQL advanced reporting, a push back to a Windows environment and away from Solaris is beginning to take place.

1. Chronos (scheduling and timekeeping app) – no virtualization. Next release will better support virtualization.
2. 18 – 24 months – Split HR and Student People Soft systems. People soft not virtualized but may be increased in the future.
    a. Tuxedo and Oracle systems split.
    b. Data growth / expansion. Double data and hardware requirements immediately.
    c. Synchronization back end system for master data management. New requirement due to systems split.
    d. Constituent data hub.
3. Virtualization in development and test environments. Not too much in production. Mainly due to storage infrastructure. Looking to move it into virtual environments for middle tier production apps.
4. All of these systems are business critical applications and need concurrent maintainability
5. Departments are currently housing their own servers and have expressed interest in having these servers hosted within the data center.
6. Campus imaging system being considered. Consolidated imaging repository would be a growth driver.
7. Bookstore/campus store may be potential collocation client.
8. Product management initiative (currently within Planview and/or XLS).

9. Master data management/Data warehousing and BI. May integrate with Health Sciences Data Warehousing or build new parallel.
10. Research/Grants proposal management systems.
11. M&A – Dixie state college and/or hosting solution for all other institutions (i.e. PeopleSoft). Separate installation of PeopleSoft most likely approach. Salt Lake community and Utah State Utah Valley University are not likely. Snow, Dixie, Southern Utah and Central Eastern more likely candidates.
12. CRM – not known what this solution will be. If not PeopleSoft would require additional hardware.
13. Additional modules will be added to PeopleSoft
14. Payroll side – twice per month paycheck printing.
15. Initiative to stop financial statements.
16. No printing – at the new data center.
17. Existing printing center would stay in place.
18. Primary printing may need to remain being backed up at Richfield.
19. Project management and facilities management group looking for collocation space.
20. ACS is currently hosting several groups:
    a. Financial and business services.
    b. UCard systems.
    c. Richard Water – Academic Affairs

Next meeting: **N/A**
Location: **N/A**

EYP MCF will rely upon these minutes of meeting as a record of the items discussed and conclusions reached. Please advise the undersigned of any additions or corrections.

Submitted by,

Robert Myers
Senior Technology Consultant

Project: **Data Center Improvements - Programming**
Client: University of Utah - Project 20109
Location: University of Utah Campus
Purpose: Hospital Applications Discovery Session
Date: May 14, 2009

Hospital Applications Discovery Session - 5/14/09			
Name	**Organization**	**Phone**	**Email**
Robert Nelson	ITS	801.587.6129	robert.nelson@hsc.utah.edu
Donald Kruppa	ITS – Ancillary	801.587.6139	donald.kruppa#hsc.utah.edu
Anne Jacob	ITS – Clinical Info Systems	801.587.6067	anne.jacob@hsc.utah.edu
Kris Lundell	ITS – Interfaces	807.587.6092	kris.lundell@hsc.utah.edu
Carrie King	ITS – Clincal Info Systems	801.587.6068	carrie.king@hsc.utah.edu
Nancy Brozeltun	ITS – Director	801.587.6187	nancy.brazeltun@hsc.utah.edu
Mike Ekstrom	OIT/ITS Manager	801.587.6086	mike.ekstrom@hsc.utah.edu
Jim Livingston	OIT/ITS	801.587.6085	Jim.livingston@utah.edu
Earl Lewis	OIT/ITS	801.581.3635	earl.lewis@utah.edu
Travis (via con call)	EMR Manager		
Steve Carter	HP/EYP - Tech. Cons. - PM	312.343.9535	scarter@hp.com
Rob Myers	HP/EYP - Tech. Cons	312.909.1567	rmyers@hp.com

Italics indicates new discussion or corrections/updates/results/response from a prior meeting

CPOE

1. CPOE - very low tolerance for down time.
2. We need to identify applications that need redundancy.
3. Design Concept
    a. Separate zone for clinical apps and very careful around those areas.
    b. Unix in one area, but citrix farm is spread.
    c. Isolate cabling, power for those systems so they do not have to meet the clinical systems rules for maintainance.
4. Growth is claims related.
5. Moving from Cerner to Epic for newer systems but not bigger.
6. Still may introduce new Cerner for greater redundancy. Cerner footprint not to change from backend perspective.
7. Shortage of non-prod servers – would like to add more non-prod servers. Would like at least 2 more 550's – Cerner.
8. Would like different security zones in data center to different groups not having to adhere to the strict health rules. The concern is what staff would make change to systems that may affect clinical systems. Would like to set up new data center that this would not be an issue.

Health Services

1. Moving toward strategy of consolodating. Will need parallel systems for bringing together.
2. Citrix Systems – 30 users per server. Large growth of Citrix systems
3. Health systems plan to grow itself and therefore growth in systems. New daybreak clinic being built, may drive decisions of adding to EPIC system, 5 year transition plan.
4. Virtualization is not supported by many vendors.

5. Talking about digitizing about more than just digital records. Significant growth

EPIc

1. Epic is the financial system, linked to ADP for up and running overnight processing. Bill 3.5 million a day.
2. Epic systems growth. Phased approach, but will require duplicate systems for some period of time. Five year transition plan.
3. Want separation of development and test systems
4. Agreeable to idea of using colo for dev/test enviroment.
5. Need to make sure space is available for systems that need to be physically next to eachother.

Space, Power, Cooling, Growth

1. Space, power and cooling is an issue, and rotate betweeen which is the problem.
2. Biggest growth rate is in EPIC. Outpatient growth has changed the most.
3. UNY PSYC is expanding there hospital, Huntsman adding 50 more beds (at 90now), ICU at main hospital is increasing (increases data). Daybreak is at planning stages – unknown growth factor
4. Visits have maitained their level

Data Warehouse

1. Backend datawarehouse is always doubling.
2. Increase space 1 to 1 of what ER create.
3. Every fin. System, EPIC, allegra, etc. feed data to data warehouse. Growth in those areas directly transfer to data warehouse.
4. Looking at taking current servers into consolodated database efforts. Currently sql servers with unconrolle growth.
5. If db is small enough we move to virtualized environment.
6. If high transaction db - host in 5 node clustered environment.
7. Best guess in no more than 10% reduction of growth of existing, but new systems increase growth.
8. Test vs build zones desired.
9. Suggest using existing DC for test which is available to staff.
10. Specific data elements for growth. 2TB of new data already committed to for growth. Since data warehouse is Epic backend environment – size of current cluster – doubling the size in next 5 years of existing cluster that serves up Epic is best esitmate.
11. IDX may be retired.

Integration

1. Well positioned for growth for next few years. MS biz talk soliution.
2. Resource limitations – has been disruptive. To current operations and future planning.

Clinical Systems

1. Capital planning, operational planning. Tie in with annual planning. Annual review of strategic plan.
2. Transcription steering committee. Current goal is to reduce the systems footprint.
3. Some systems may need to stay close to the hospital.
4. Proximity issues will still be an issue with some medical systems.

5. Netcom issues
6. Dial tone access – model access still required.

PBX

1. Some systems need to be in close to phone switch, may not be able to move,
2. If new DC has an appropriate PBX, etc. then these type of systems can be moved. Will need POTS line, etc. to make sure we meet requirments of some systems.

Distance Limitation Servers

1. The new west pavilion of Hospital there is a space that can support approx 20 racks.
2. Plan is to move systems that must be located at the hospital will be consolodated into this new room.
3. Thinking of using new data center for backup of these systems – one or two servers.

Ancillary Systems

1. Currently 93 systems. Cannot anticipate growth
2. Would like to focus on consolodation to reduce footprint.
3. Systems currently located in hopital, and data center. Would prefer to have all systems located at NDC. 15-20 systems.
4. Model for at least a cabinet for now.
5. Would like to separate dev and prod developments.
6. Would like to isolate business and clinical systems in different areas.

Financial Systems

1. Business continuity and bandwidth is the concern.
2. 0 tolerance for delays or downtime process.
3. Future growth areas would be webservices for patient services. Portal services, Telemedicine – registering patientss.

Imaging

1. PACS growing consistantly.
2. Genome is handling at ARUP not at the hospital itself.
3. 5 yr ago 15 tb….now at a petabyte.

Next meeting: **N/A**
Location: **N/A**

EYP MCF will rely upon these minutes of meeting as a record of the items discussed and conclusions reached. Please advise the undersigned of any additions or corrections.

Submitted by,

Robert Myers
Senior Technology Consultant

Project: **Data Center Improvements - Programming**
Client: University of Utah - Project 20109
Location: University of Utah Campus
Purpose: Campus-wide Services Discovery Session
Date: May 13, 2009

Campus-wide Services Discovery Session - 5/14/09			
Name	**Organization**	**Phone**	**Email**
Brad Grow	Media Solutions	801.587.7664	brad.frow@utah.edu
Shellie Eide	Media Solutions	801.585.9838	seide@media.uatah.edu
Earl Lewis	OIT	801.581.3635	earl.lewis@utah.edu.edu
Scott Allen	UEN	801.581.5382	scott.allen@uen.org
John Desha	UEN	801.581.4778	desha@uen.org
Paula Millington	Media Solutions	801.581.3032	paula.millington@utah.edu
Jessica Stokes	OIT/ACS	801.585.0688	jessica.stokes@utah.edu
Caprice Post	OIT/ITS	801.585.5404	caprice @utah.edu
Kevin Taylor	OIT	801.585.3314	kevin.taylor@utah.edu
Steve Carter	HP/EYP - Tech. Cons. - PM	312.343.9535	scarter@hp.com
Rob Myers	HP/EYP - Tech. Cons	312.909.1567	rmyers@hp.com

Italics indicates new discussion or corrections/updates/results/response from a prior meeting

OIT/ACS

1. Concerns with availability, how would it communicate with richfield, redundant power.
2. Voice systems must be able to be continued to operate. Systems are sitting in various data centers must be highly available. Unified messaging. May be moving to VOIP Call Managers
3. Intend to run on virtualized systems.

UEN

1. 3 more year contract with blackboard Vista. After that not sure on system.
2. Bringing on line Moodle from Blackboard in to replace virtual high school. Will expand the existing hardware.
3. Concerned about uptime and reliability and many schools depend on enterprise applications.
4. 17 full racks. Trend moving towards blade servers. 2-3 KW per chassis.
5. Migrating from tape to disk backups. Will be running trays of disk.
6. Working on getting blackboard fully redundant.
7. Want to run in multiple locations. And coming in geographically node with failover.
8. Running out of cooling.
9. 3 times power issues would like reliable power and properly maintained equipment. Goverance and enforcement of policies.

Media Solutions

1. Works with student and faculty data - mostly course data. No sensitive data.
2. Large amount of faculty data for running reports. Must be available.
3. Higher reliability needed for obtaining large grants.
4. 850 web sites on the home site.
5. Other web sites being down reflect on the Universities image.

6. Portal, content and access.
7. Collaboration application – Vingette. Unite. Rolling out for researchers first.
8. Retention, promotion and tenure for faculty.

Next meeting: **N/A**
Location: **N/A**

EYP MCF will rely upon these minutes of meeting as a record of the items discussed and conclusions reached. Please advise the undersigned of any additions or corrections.

Submitted by,

Robert Myers
Senior Technology Consultant

Project: **Data Center Improvements - Programming**
Client: University of Utah - Project 20109
Location: University of Utah Campus
Purpose: Disaster Recovery / Business Continuity Discovery Session
Date: May 14, 2009

Disaster Recovery / Business Continuity Discovery Session - 5/14/09			
Name	**Organization**	**Phone**	**Email**
Scott Goodel	ITS	8001.587.6126	scott.goodell@hsc.utah.edu
Guy Adams	CHPC	801.554.0125	guy.adams@utah.edu
Earl Lewis	OIT	801.581.3635	earl.lewis@utah.edu.edu
Steve Corbeto	OIT/CI	801.585.9464	steve.corbeto@utah.edu
Brent Elieson	OIT/ITS	801.587.6030	brent.elieson@hsc.utah.edu
Caprice Post	OIT/ITS	801.585.5404	caprice @utah.edu
Bryan Peterson	ITS	801.587.6095	bryan.peterson@hsc.utah.edu
Christian Shank	ITS	801.213.3315	christian.shank@hsc.utah.edu
Jeff Hadden	ITS	801.587.6041	jhadden@hsc.eutah.edu
Bryan Peterson	UEN	801.585.7789	bryan@uen.org
Greg Nance	ITS	801.587.6108	greg.nance@utah.edu
Shannon Thayn	ITS	801.587.6060	shannon.thayn@hsc.utah.edu
Mike Ekstrom	OIT/ITS	801.587.6086	mike.ekstrom@utah.edu
Tim Richardson	ACS	801.585.1152	tim.richardson@utah.edu
Glen Cameron	ITS	801.581.4290	gcameron@acs.utah.edu
Bryan Morris	OIT	801.585.7789	bryan.morris@utah.edu
Steve Carter	HP/EYP - Tech. Cons. - PM	312.343.9535	scarter@hp.com
Rob Myers	HP/EYP - Tech. Cons	312.909.1567	rmyers@hp.com

Italics indicates new discussion or corrections/updates/results/response from a prior meeting

OIT/ITS

1. Move all production into new data center.
2. Richfield as DR for top tier applications.
3. UEN prescince for the network connectivity capability from new data center to Richfield
4. Need connectivity that bypasses campus in case of campus outages.
5. Run book issues and operations procedures.
6. Plan for MAN is to hit two sites on campus –EBC and Park building, new data center, Level 3, World telecom association.
7. MAN ring would cross the major fault line.
8. UTA conduit – slack in the fiber – working to configure slack to allow for some movement of fiber connectivity. Wireless link will be emergency backup due to crossing fault line.
9. MAN might be expanded to state wide education network, driving south down I15 to Richfield.
10. Payroll, financial systems, registration. Hospital – core services, EMR, financial services going to Richfield for DR.
11. Emergency preparedness from a patient perspective existing for the Hospital perspective. Tier 1 applications located at Richfield. Now that Hospital is part of CPOE, IT uptime is critical
12. Physicians billing and radiology - 72 hour RTO. Spare server copied hourly at Richfield.

13. Radiology is currently evaluating RTO. Radiology would most likelu utilize Richfield as well.
14. Critical radiology system PACS is filmless.
15. PACS requires high bandwidth. PACS moving to ASP model but does host the data. Data will be campus and at ASP backup locations.
16. If moved to datacenter should have isolated wavelength for PACS.
17. PACS image sizes are increasing exponentially.
18. Hardware owned my PACS providor – may need to be treated as Colo.
19. Need to have the physical and logical network redundancy to handle most of the scenarios
20. DNS, AD have presence in Richfield. Need to have web hosting services in Richfield.
21. Ring must have multiple paths.
22. Need ability for clustering between new data center and other data centers to move apps back and forth; need the bandwidth to accommodate this need.
23. Need 10 gig connectivity. Near active/ative replication on databases.
24. UNIX dbase would be hosted.
25. Windows servers and sql servers.
26. Citrix servers need for user access.
27. Tape library needed.
28. Block based data replication – not available over 1 GB disk systems.
29. Upgrade to enterprise class storage required at Richfield to meet high speed replication needs.
30. EPIC is requesting to have replications through dbase.
31. NetAPP – likely to replicate to Richfield from DTC.
32. Better infrastructure for the fiber channel – max current is 1 Gig limitation.
33. Need to make sure that there will be adequate dedicated capabilities and capacities to handle database replication needs.
34. Ability to expand beyond what is considered high capacity today for unknown requirements in the future.
35. EPIC, Cerner – have tools to replicate at the transaction level. May need to replicate in the future at the block level with SAN hardware.
36. Cache and EPIC does not utilize Oracle replications.
37. Operations control and change control needs to be address at Richfield as well as the new data center. Processes needed.

CHPC

1. Research – not a significant DR plan in place.
2. Grant funding usually only enough to cover primary equipment
3. HPC – major recovery would be to ask for time from other institutions.
4. Network access needed to access other external resources.
5. CHPC – silo from IBM. Quarterly backups offered to major researchers.
6. Compliance may dicatate requirments such as data access, open access.
7. Library beginning to look at as role as data currator for grant based data.

UEN

1. DNS is replicated throughout the state.
2. Course management system using oracle replication and is successful.
3. Future plan is to have application nodes at Richfield. Should be in place by fall – manual fail over.
4. Mirroring of SAN synchronous replication downtown and tertiary in Richfield asynchronous.
5. Looking to replicate storage fabrics between EBC and downtown datacener riding on metro ring.
6. Going to virtual tape library from tape, not sure if in EBC or in new data center.

ACS

1. Support Enterprise portal, financials and students systems.
2. Moving DR site to Richfield next Wednesday – Oracle Dataguard.
3. Student systems DR from Richfield now, would like to be able to do all services through Richfield as backup.
4. Future applications Crystal Reports and all other major systems.
5. Credit Card processing is an issue over current data pipes.
6. 4540 Thumpers to ship data between data centers.

Next meeting: **N/A**
Location: **N/A**

EYP MCF will rely upon these minutes of meeting as a record of the items discussed and conclusions reached. Please advise the undersigned of any additions or corrections.

Submitted by,

Robert Myers
Senior Technology Consultant

Project: **Data Center Improvements - Programming**
Client: University of Utah - Project 20109
Location: University of Utah Campus
Purpose: High Performance Computing (HPC) Discovery Session
Date: May 15, 2009

High Performance Computing (HPC) Discovery Session - 5/15/09			
Name	**Organization**	**Phone**	**Email**
Julio Facelli	CHPC	801.585.3791	julio.facelli@utah.edu
Joe Breen	CHPC	801.585.1013	joe.breen@utah.edu
Guy Adams	CHPC	801.554.0125	guy.adams@utah.edu
Steve Corbeto	OIT/CI	801.585.9464	steve.corbeto@utah.edu
Martin Euma	CHPC	801.652.3836	m.cuma@utah.edu
Philip J. Smith	CHPC	801.585.9464	philip.smith@utah.edu
Thomas Cheathem	CHPC	801.587.9652	tec@utah.edu
Charles Prawdzik, Jr	HP/EYP - Architect - PM	310.864.3268	cprawdzik@hp.com
Steve Carter	HP/EYP - Tech. Cons. - PM	312.343.9535	scarter@hp.com
Rob Myers	HP/EYP - Tech. Cons.	312.909.1567	rmyers@hp.com

Italics indicates new discussion or corrections/updates/results/response from a prior meeting

1. 750 KVA at Komas data center - Maxed out
2. Estimate 1-1.5 MW – of HPC Clean power (15% of total power on UPS) will suffice for immediate needs and near future. Projecting future needs is not possible.
3. 250 KVA UPS estimated need for critical storage
4. Maybe moving to 4-6 MW in the future.
5. Would like connections for chilled water to the racks – loops installed for chilled water. Not installation of chilled water plant.
6. Four foot high raised floor desired for being able to add chilled water cooling, etc.
7. Interested in Stanford design – ambient air cooling systems.
8. Floating point units – generate significant loading for power and cooling. 120 degrees exhaust temperatures.
9. Hot aisle air containment may be necessary
10. Dock – storage, build area. 750 SF. 6 KW of power
11. 3 Phase for the high density PDU's.
12. 8 cubes shared office space requested at new datacenter
13. Conference room to accommodate 20 people requested.
14. HPC install their own equipment. Needs to be separated from Enterprise system. Separate security access.
15. 480 V – flexible power distribution.
16. DC plant for the colo plant and for network.
17. 20 KW per cabinet average density for power.
18. 9 Racks for last cluster – 3000 SF
19. 5000 SF day 0 – Assumptions non large - Tier 2 systems.
20. An unexpected large research project would require separate HPC facility expansion project.
21. Monitoring required in new datacenter – airflow, power and humidity for every cabinet, water on the main floor. UPS to rack level monitoring. Data points and trending is important.
22. Mechanical/Electrical Efficiency/Operations tracking – theoretical vs. actual.

23. Integrate systems level monitoring from computing platforms. Building management system assumed to be part of this new data center facility.

Next meeting: **N/A**
Location: **N/A**

EYP MCF will rely upon these minutes of meeting as a record of the items discussed and conclusions reached. Please advise the undersigned of any additions or corrections.

Submitted by,

Robert Myers
Senior Technology Consultant

Project: **Data Center Improvements - Programming**
Client: University of Utah - Project 20109
Location: University of Utah Campus
Purpose: Collocation Requirements Discovery Session
Date: May 15, 2009

Collocation Requirments Discovery Session - 5/15/09			
Name	**Organization**	**Phone**	**Email**
Andrew Reich	OIT	801.587.0902	andrew.reich@utah.edu
Earl Lewis	OIT	801.581.3635	earl.lewis@utah.edu
Steve Corbato	CI Strategy	801.585.9464	steve.corbato@utah.edu
Joe Breen	Center for HPC	801.550.9172	joe.breen@chpc.utah.edu
Bren Elieson	ITS	801.587.6041	jhadden@hsc.utah.edu
Bryan Morris	ITS	801.585.9229	bryan.morris@utah.edu
Jim Livingston	OIT/ITS	801.587.6085	jim.livingston@utah.edu
Bill Billingsley	Design and Construction	801.585.0073	bill.billingsley@fm.utah.edu
Charles Prawdzik, Jr	HP/EYP - Architect - PM	310.864.3268	cprawdzik@hp.com
Steve Carter	HP/EYP - Tech. Cons. - PM	312.343.9535	scarter@hp.com
Rob Myers	HP/EYP - Tech. Cons	312.909.1567	rmyers@hp.com

Italics indicates new discussion or corrections/updates/results/response from a prior meeting

1. Need to decide if colleges should be in enterprise area or in colo
2. University Enterprise Data Center Services will be the label of the core.
3. Network – clinical care, research, and lower campus
4. Recommendation of evaluating on an application basis as opposed to department basis.
5. Will need a process to decide whether apps are enterprise or colo slated.
6. Maybe all depts will go to colo until there is a business case to bring into enterprise and approval.
7. Ucard, bookstore, GIS, Develompment Office may be possible candidates for collocation.
8. Possible Colleges – Law, Fine Arts, Architecture, Business, Behavioral Sciences, School of Medicine, Colleg of Nursing, Colleg of Pharmacy, College of Health, Mines and Engineering, Education, Social Work.
9. Intent is to keep printing out of the DC space
10. Project Management Group and Facilities Management Group looking for data center space.
11. Possible University Departments - Facilities – Campus Design Construction, Planning, Operations, Student Systems, Financial Imagining - FORTIS
12. ACS is currently hosting several groups:
    a. Financial and business services.
    b. UCard systems.
    c. Richard Water – Academic Affairs
13. Possible entinities outside of U of U
    a. ARUP labs
    b. Intermountain Healthcare
    c. State board of regents

d. Utah State and Southern Utah for HPC area. Primary focus on Utah State.
e. May be DR site for Southern Utah
f. DTS - State Dept of Technical Services
g. State Dept of Health – possible medical collaboration
h. American Geological Institute
i. HHMI
j. Brain institute
k. USTAR
l. EGI Energy and Geophysics
m. SCI - Scientific Computing Institue.
n. Huntsman Cancer (Tony Morillo)
o. Clinical Care, Research and Lower Campus
p. Conflict of interest – Erica systems

Actions:

EYP MCF

1. Create a questionairre for possible collocation candidates to measure interest and provide to U of U. *(Issued 05/21/2009)*

U of U

1. Interview potential candidates of collocation and provide details back to EYP.

Next meeting: **N/A**
Location: **N/A**

EYP MCF will rely upon these minutes of meeting as a record of the items discussed and conclusions reached. Please advise the undersigned of any additions or corrections.

Submitted by,

Robert Myers
Senior Technology Consultant