AE Initiative Business Case – Data Center Aggregation

Initiative Sponsorship and Ownership

Project Name:	***IT – Data Center Aggregation***
Project Summary:	The objective of this project is to analyze and identify cost savings and efficiencies by assessing the existing server/data center infrastructure and making recommendations for the future state. Areas of examination and data collection include: utility costs, labor efficiencies, effective best data center practices, and space re-utilization costs, with the goal of improving risk management and the security of sensitive data through the co-location and virtualization of servers and data centers across campus. The goal of this project is to develop a new model of server and data center structure to serve the needs of the University's academic, research, and administrative communities, leveraging industry-leading practices for server administration, virtualization, and management to save costs, improve service levels, and minimize data security risks.
Business Unit(s):	Vice Chancellor for Administration – Administrative Excellence
Business Process Owner(s):	CIO/proposed organization under the CIO
Preliminary Cost Estimate:	Variable ~$200-600K annual cost based on level of virtualization and co-location; initial capital investment (if necessary) to be determined by implementation team
Preliminary Savings Estimate:	~$6.8M over 5 years
Proposed Go-Live Date:	Estimated to be 30-45 weeks following initiation of implementation phase

Business Need or Opportunity

Current State Summary

Currently, UW-Madison has no central campus level approach for providing data center services. Therefore, the server/data center model on campus is inefficient, resulting in duplication and overspending in areas including: hardware purchases, utility costs for power and cooling, labor, and facilities. There is little coordination, consistency or communication in putting up a server on campus today and no incentives in place to consider a central facility or hosting service.

As it stands, there are a multitude of server rooms across campus that provide a wide range of hosting services. The team identified 97 facilities through its March 2012 survey of data center administrators:

- Approximately 52,000 sq. ft. of identified hosting space; approximately 55 percent used for research purposes
- 60 percent of facilities host virtualized servers
- 25 percent of facilities are 500 square feet or larger; 16 percent are 100 square feet or smaller
- 50 percent of facilities cannot freely share their facilities with others due to restrictions on how their funding can be used or particular security policies they must follow
- 53 percent have lockable racks; 34 percent have card access systems
- 76 percent of facilities have a dedicated cooling system; 30 percent have a backup cooling system

The current primary driver for developing server room space leans towards expediency and convenience with little emphasis on reliability, energy efficiency, or security. The result is that most of these data centers are not managed according to industry norms and are inefficient. None are managed as a campus resource. Some aren't properly sized. Almost all were designed before the advent of virtual computing.

For some departments with access to facilities and services, recent virtualization efforts have reduced the need for expansion into additional spaces. Meanwhile, other departments who need server room services but don't have access to them must beg, borrow, or lease space, or hope that their department has enough money to modify an existing space appropriately. Ultimately, to minimize cost and space, some end up utilizing suboptimal spaces for server hosting. While the biggest return on investment is in encouraging the virtualization trend, smaller departments don't have the resources to take advantage of this technology effectively. Departments with resources for virtualization are building similar infrastructures across campus at the departmental level. The result is "haves" and "have-nots" in an inefficient and wasteful environment.

Opportunity

The opportunity exists today to encourage best data center practices using incentives that will result in significant cost savings for central campus, individual schools, colleges, and administrative units through the reorganization and consolidation of server placements and solutions including increased virtualization, aggregation, and outsourcing.

The primary driver of this project is to be able to recover cost savings from co-location and virtualization. However, other significant considerations include the growing requirements to provide safe, secure storage for sensitive data. There exists an increasing need to safely store data, and the current server/data center environment puts the institution at risk for data loss and the resulting financial and reputational costs.

The mission of the University may not include building and managing a high-end data center today but rather to leverage available options and assets in managing risk, securing sensitive data and responsibly managing costs. The University is dependent on a wide range of IT resources to carry out its mission of teaching and research. This project provides the opportunity to inventory and assess the fitness of the existing infrastructure in support of that mission and recommend changes as needed.

Proposed Solution Description

Scope

Methodology

- Reviewed data from Huron's phase I work
- Built glossary of common terms
- Performed a scan of campus network segments to provide an objective baseline on current server infrastructure
    - Vetted scan results with campus units to clarify estimates
- Held a listening session with data center administrators; distributed a follow-up survey which resulted in a detailed inventory of 97 facilities on campus
- Collected power usage effectiveness (PUE) data to gauge energy efficiency of various on-campus data centers/server rooms
- Assessed existing virtualization efforts on campus from a methodological and financial perspective (e.g. CALS, DoIT)
- Conducted site visits to on-campus and off-campus facilities (TDS)
- Drafted a high-level service model as an example of a future offering to campus

Recommendations

The data center aggregation team recommends:

- A reduction in the number of campus facilities dedicated to housing servers
- The creation of a campus shared data center service provider

The team recommends that all data center facilities under consideration for continued operation have minimal levels of energy efficiency, security, and projected uptime. In addition, any data center facilities used in the provision of campus-level services must exceed those minimums depending on the type of services to be provided.

The team proposes that the minimum level of security for any data center must include a mechanism to control access to only those individuals needing it, a mechanism to record that access, and an expedient mechanism to remove access. For that reason, we feel that any facility to be used as a data center be equipped with card reader access tied into campus security. A higher level of access control must be considered in data centers participating in campus level services since restricted data may be stored there as a result. A security baseline for campus services should be determined to be as inclusive as possible with known regulations.

The proposed campus service will manage server hosting facilities as a shared resource across campus. It will build a set of co-location and virtualization services on top of the facilities services. The central service would provide a one-stop shop for data center services on campus as follows:

Central Service Summary

- Reports through the CIO
- Guided by the equivalent to a steering committee representing major stakeholders
- Pools expertise already existing on campus to staff the unit
- Uses a service layer approach to provide managed data centers as a campus service
- Uses campus based and private service providers as appropriate
- Provides consistent services across campus to all departments
- Creates a foundation of services that departments could use directly or use to build a suite of services targeted to a particular customer base (e.g. administrative computing; research computing; etc.)
- Leverages savings from reducing physical infrastructure to fund central service

Foundational (Common Good) Services Provided

- Facilities management
- Server co-location service
- Virtual server service
- Consultation service
- Operational coordination (policy/procedure surrounding access, accountability, governance, capacity planning)

Goals

- Promote and incentivize best practices
- Provide a consistent level of service to all departments that is easy to use and flexible
- Maintain an inventory of server hosting facilities on campus
- Manage selected server hosting facilities consistently and according to industry best practices
- Take advantage of private service providers when it makes sense
- Eliminate substandard server rooms
- Match hosting facilities with service needs
- Eliminate duplicative infrastructures at the departmental level while maintaining service levels
- Encourage virtualization
- Enable service continuity throughout transition

Overview - Managed Server Room Services

The foundational services provided must be underpinned by a set of server hosting spaces that are matched to the service levels and features required by customers. The central service will be responsible for identifying and managing these spaces on behalf of campus and for using them to provide the foundational services. Based on industry trends, the team expects that the campus' need for server hosting spaces will be best met by a combination of third-party and on-campus server room services sources. Virtualization and use of private service providers will slow or eliminate the near-term need for new campus facilities. The on-campus spaces must be managed using standard techniques and cost models for all spaces. All server hosting spaces within the same service layer should provide consistent services.

The team recommends that implementation work finalize an inventory of campus spaces, assign values to these spaces, categorize customer's needs, and match a set of on-campus and private services to these needs. The central service will manage these spaces and service contracts with private providers on behalf of campus.

Overview - Co-location Service
The team recommends that the server room spaces that are not going to be used to provide the central service be abandoned as the existing servers in these spaces reach the end of their lifecycles. Customers that need physical servers would use space provided by the central service as described above. The services provided by the central service would range from simply providing rack space to managing servers on the customer's behalf in a highly secure environment. Implementation work includes determining the features of the co-location service.

Overview - Virtual Server Service
The biggest "bang for the buck" is server virtualization. The team recommends that the central service provide a virtual server service including disk storage. This service would range from automatic virtual server creation via web interface to physical to virtual conversion services. The virtual servers should be hosted on campus-provided infrastructure or via a cloud service (private, public, commercial, or hybrid) as appropriate. The server hosting arrangements will be managed by the central service and transparent to the user. The potential service offerings range from basic virtual servers that users can customize to their needs to fully managed servers that meet security requirements such as HIPAA, and PCI-DSS. The feedback the team has received indicates that a successful virtual server service must be flexible, agile, and cost less than standing up an equivalent service at the departmental level. Implementation work includes determining the features of the service, platforms for providing the service, interfaces for requesting the service, cost structure, and return-on-investment breaks.

Assuming that several relatively large enterprise quality data centers will continue to operate and that virtualization services be made available, a logical step would be to consider the creation of a campus cloud via an infrastructure as a service model (IaaS). A private cloud has the advantages of cloud computing (i.e. redundancy, dynamic allocation of resources, etc.) without the potential regulatory, legal and connectivity issues related to the use of outside vendors. In addition, depending on the pricing model used by a vendor, a local private solution may be less expensive. Several hardware vendors are offering or are in the processes of developing IaaS platforms and the decision to develop such a service would at least partly be dependent on that market.

Overview - Consultation Service
Deciding which service or services to use assumes that customers have a technical understanding of information technology. Many departments can't afford high-level staff with this knowledge. Our listening session results indicate that many customers want to have access to a group of experts who can help customers select the best options for them among the virtualization and co-location service offerings. The team recommends that the central service includes consultation services. The consulting services provided would range from assisting with the selection of options to complete analysis and implementation consulting surrounding best practices and physical to virtual conversions. Implementation work includes defining the consultation services to be offered by the central service.

<u>Proposed Policy Framework</u>
The project team recommends the following framework be considered in development of policy language:

- To whom does it apply

    - Individuals who are making institutional purchases of physical or virtual servers
    - Data center facility managers
- Process
    - Over time, expectation that purchasers/process owners/facility managers will consult with/turn to the central service provider to ensure the most efficient and effective use of resources and determine the right path to achieving their desired service level
- Compliance/oversight
    - Responsibility for oversight needs to be distributed and occur at multiple levels (e.g. VCA, business process owner, Dean's Office/Department leadership, etc.), rather than falling solely on the central service provider
    - "Oversight" includes:
        - Pre-approval of exceptions

Additional policy will be needed to define when a new data center is created, when an existing data center is modified or enhanced, and when a data center should be decommissioned. Data center policies should also be consistent across all campus facilities. This can be defined through the data center governance and advisory structure and enforced through the CIO's office. Policies for access controls, auditing, space utilization and planning, networking, facilities management, inventory/asset management, and others will need to be defined.

Implementation Framework

The team recommends that three implementation teams are formed to: define the organization/governance of the central service, define the set of services to be offered by the central service provider, and identify the set of hosting platforms where and how these services will be provided.

- Central Service Organization/Governance Team
    - Work closely with CIO
    - Develop funding model
    - Develop operational model
    - Develop staffing model
    - Develop attendant policies
- Services Team
    - Interview customers to determine what services must be provided
    - Set service level expectations
    - Define server hosting facility attributes/requirements
- Data Center Facilities Team
    - Complete campus data center inventory
    - Identify spaces suitable for providing co-location and virtualization services
    - Determine level of investment needed to bring existing spaces to required campus standards
    - Evaluate cost/benefit of alternative hosting options
    - Identify services to be provided by off-campus hosting vendors

Proposed Milestones and Timing

Milestone	Timing / Date
• Identify individuals to be on implementation teams and team member responsibilities/roles • Schedule initial team meeting	3-4 weeks
• Develop and finalize detailed implementation plan	4-6 weeks
• Collect customer service level requirements and build service model	10-12 weeks

• Complete data center inventory; visit and assess prioritized spaces; determine level of investment required	12-15 weeks
• Assess off-campus hosting options; vendor engagement	8 weeks
• Determine operational and staffing model	8 weeks
• Determine funding model	4-6 weeks
• Develop communication plan and change management plan	4-6 weeks
• Develop policy/procedure for purchase of servers, menu of services documentation	4-6 weeks
• "Go Live" of Proposed Solution	30-45 weeks after initiation

Alignment with Strategy

This solution will align with many of the current UW-Madison strategic priorities and initiatives, as well as those strategies identified by the institution's functional areas.

Applicable Strategic Objectives	Alignment with Strategy
University Strategy – For Wisconsin and the World, Campus Strategic Framework (2009-2014)	
• Be responsible stewards of our resources • Align resources with priorities • Make our administration more effective, efficient, and flexible	• Reduction of physical servers and inefficient facilities will decrease spend on utilities and hardware and reduce burden on distributed IT staff, allowing for additional focus and resources to be directed towards other initiatives • A tiered service model allows for flexibility in delivery and ensures that customers aren't paying for services they don't require
Institutional IT Strategy – IT Strategic Plan (2010-2014)	
• Leverage enterprise infrastructure and avoid unnecessary replication of infrastructure and services • Maximize transparency across the campus • Work toward green computing strategies	• Positioning of central service as single point of contact utilizes enterprise infrastructure • Channeling the acquisition of a majority of servers and related items through enterprise sales channels reduces redundant purchases and increases the transparency of the server infrastructure across campus • Standardization of tiered services results in more consistent and less costly IT support • Server virtualization and consolidation is a core component of a green IT infrastructure through reducing energy consumption • Alignment with UW-Madison Advanced Computing Infrastructure (MACI)
Functional Area Strategy – VCA Strategic Plan (2009-2014)	
• **Resource Stewardship**: Improve services and clearly demonstrate to campus customers and the public that resources are used responsibly by: • Improving process efficiencies in order to enhance services and responsiveness to campus customers as well as identify cost savings and improve the institution's financial performance. • Sharing services across VCA units and with VCA partners to increase collaboration, reduce redundancy and	• Simplifying the virtualization and server acquisition process leads to greater efficiencies across campus • Reduced time spent by customers/providers to efficiently obtain and deliver services will result in cost savings • Reduced time spent researching servers, VMware tools, and other components will result in cost savings • Maximizing the use of underutilized and efficient facilities will reduce unnecessary investment and duplication • Leveraging collective UW expertise in a central service will increase collaboration and expand the sphere of influence

Applicable Strategic Objectives	Alignment with Strategy
duplication, and free up resources for reallocation	

Anticipated Benefits

Benefit Categories	Description
Improved Service	• A central resource for data center/virtualization needs will support divisional IT staff with the expertise to assess, recommend and deliver the optimal solution. The central resource will result in higher availability of support staff, increased speed of service requests, and rapid innovation
Reduced Costs	• Gross savings of ~$6.8M over five years following implementation are outlined in the financial model section and appendix
Enhanced Engagement with IT Community	• The central resource will work closely with division level IT staff, developing trust through engagement. The steering committee will also provide a vehicle for collaboration and engagement across campus
Staff Time Freed for Other Priorities	• Divisional/departmental IT staff will be able to trust the expertise provided in the central service and will be able to focus on their division's strategic priorities, including research and instruction
Mitigated Compliance Risk	• The central resource staff will develop a wide range of expertise to ensure compliance for campus data needs (e.g. HIPAA, FERPA, evolving regulatory requirements of granting agencies)

Customer Readiness

The diverse nature of the IT and user community at UW-Madison was reflected in the team's stakeholder engagements, which included a listening session and survey of data center administrators, a server quantification process across campus, and a meeting with MTAG. While a number of concerns were raised about a central service model, an appetite among certain stakeholders to relinquish varying degrees of the server administration process provided the impetus for our recommendation. The following are the key findings from the team's stakeholder engagements:

Facilities

- Through an extensive survey, the team identified 97 server room/data center facilities across campus. Of these facilities, eight percent are offering a paid hosting service, and eighteen percent are interested in doing so:
    - Indicates that a number of facilities might have the resources and infrastructure to support the needs of customers outside of their immediate domain
    - Demonstrates potential interest among a separate set of administrators to offer their services to others
    - May indicate a willingness to collaborate with a managed campus service
    - **Concern:** Administrators are not thinking in terms of a shared central service, but rather a revenue generator for "their facilities"
    - **Concern:** Departments/units have invested their own funds to build facilities, may be reluctant to relinquish unused space for campus needs
    - **Mitigation:** Explore incentives for collaboration (e.g. facility upgrades and maintenance); develop effective transition management plan to aid cultural change
- Based on the survey responses, the team reviewed high-level key characteristics for potential hosting sites and identified types of facilities that may be candidates for the high-efficiency data centers that the campus could seek to maintain under the proposed model:
    - Facilities include DoIT's in the Computer Science building, the School of Medicine's in the Medical Foundation Centennial Building, SSCC's in the Social Sciences building, the Graduate School's in Henry

Mall, Education's in the Education building, Chemistry's in the Chemistry building, and space in the basement of Ingraham Hall
    - The majority of these facilities have raised floors, indicate available space, and are willing to host
    - **Concern:** Facilities may warrant additional investment
    - **Concern:** Data was limited to respondents that provided information; not all-encompassing
    - **Mitigation:** Complete data gathering effort and conduct site visits/interviews with responding administrators to assess facility readiness/determine most cost-effective approach
- 40 percent of survey respondents indicate that there are no virtualized servers within their facilities
    - While not all servers are suitable for virtualization, this number indicates that there may be an opportunity for increased virtualization
- Across all facilities surveyed, 40 percent of respondents reported that their data centers were fully utilized. However, for the other 60 percent of respondents, an average of 33 percent of space was reported as not utilized
    - Presents both an opportunity for growth within particular facilities and consolidation of underutilized and inefficient facilities
    - **Concern:** Space is self-reported and may not be useable for server hosting
    - **Mitigation:** Formal walk-through evaluation could identify useable vs. unusable spaces

Servers

- Through the administrator listening session, participants relayed both support and concern for end users that would be asked to consider co-locating or virtualizing their servers through a central service model. Overall, the team's position is that end users exhibit varying degrees of readiness:
    - Some end users are already virtualizing and feel they can do it better themselves
    - Others want nothing to do with the physical boxes and hardware and would be early adopters
    - Others still want to manage the operating system and patches
    - Other only want to manage the application
    - Others are completely resistant and will need significant convincing
    - **Concern:** Many end users are apprehensive because of the perceived lack of customer focus, limited service tiers, and prohibitive cost model associated with DoIT's current offerings
    - **Mitigation:** Relay differences in the new model through an effective communication plan; rely upon positive word of mouth through effective engagement with early adopters

Summary

There are a number of facilities that demonstrate both efficiency and available capacity and are therefore likely candidates for hosting and partnering with a central service. Other facilities are inefficient and underutilized. The majority falls somewhere in the middle. Administrators within all three groups have demonstrated varying levels of appetite for the central service model. For those whose facilities are targeted for eventual closure, there will be resistance to protect investment and desired service levels. A strategic approach will need to be developed in order to target inefficient operations while continuing to meet the needs of end users.

There are a number of end users that would immediately consider partnering with a central service to administer their server hardware. The centralized virtualization offerings will be of particular interest to this group. Others will demonstrate significant resistance due to concerns over customer service and cost. A customer engagement initiative will need to occur to further define desired service tiers and ensure needs will be met at various levels. The change management and communication plans will emphasize the differences between the new solution and current model.

Stakeholders Impacted

The team acknowledges that the introduction of the recommended solutions will affect various stakeholder groups. These impacts are described below.

Internal Stakeholder Group	Impact
CIO & Staff Deployed to the Central Resource	• Staff identified to be part of the central resource group will be re-deployed as the consultants/experts in data storage and processing
Computer Support Staff/Server Administrators	• Will now work cooperatively with the central provider to determine optimal solutions
Computer End-Users	• Solution will result in improved services and solutions for their data needs
Deans/Directors/Departmental Administrators	• Will provide some level of oversight for compliance • Will be involved in the development of metrics • Will be involved in the development of exception policy • Will support, encourage and ensure that the central resource is used by their division/department
P-Card Administration	• Will be involved in the purchasing framework for servers
UW-Madison Purchasing	• Will oversee standard procedure development for the purchase of servers and the exception process • Will be involved in policy development/enforcement
Business Offices and IT Groups	• Will provide oversight for compliance – use of central service and server purchasing • Education/training
Division of Information Technology (DoIT)	• Will have some staff re-deployed to be a part of the central service team
Steering Committee/Executive Leadership	• Will provide oversight for the development of the implementation teams and compliance

External Stakeholder Group	Impact
Vendors (Servers, Virtualization, Hosting Services)	• Will be involved in negotiations with UW-Madison • Will be responsible for development of hardware discount options for central purchasing

Impact on Other Initiatives

There are several initiatives underway which may have synergies with the proposed solution contained in this document; however, some initiatives may also be challenged and/or enhanced by this proposed solution.

Initiative	Impact
MACI - The University of Wisconsin-Madison Advanced Computing Infrastructure	• The implementation teams will need to work closely (or combine efforts) with the proposed MACI solution to support research computing needs • Ideally, the central service would develop the expertise to support administrative and research computing needs
We-Trade (FP&M)	• Creative solutions for the space available in divisional data centers might be sought
Other Administrative Excellence Teams: • Resource Allocation team	• Consolidated server/data storage options could lead to the reallocation of resources to serve campus' data storage/processing needs

Initiative	Impact
• Policy team	• New policies will be developed, assessed, and maintained to support this initiative
UW Purchasing Initiatives	• Current contracting and vendor initiatives/approaches may change and/or be enhanced as a result of this proposed solution
IT Initiatives	• There may be some impact on any current departmental, AIMS, and DoIT projects related to data management
Sustainability	• May impact existing sustainability initiatives related to virtual or consolidated servers

Project Success Factors

Change Management Plan

Overview

The communication plan will need to concisely frame a very complex and wide-ranging subject, inform key stakeholders throughout the assessment and project implementation phases, and provide benchmarks so the larger campus community can readily assess progress towards transition goals

- Stakeholders to be targeted include:
    - UW-Madison leadership
    - Current campus data center managers, server administrators, and IT service providers
    - Deans and department chairs supporting distributed IT operations
    - Governance groups

The domains of data centers and of virtualization are large and complex. There are a number of formal and informal service and business models in use on campus, each with its own benefits and risks. To fully grasp the work of this team and the impact of current and future changes, stakeholders will need to understand some of the cultural and technical complexity in play. Therefore, the communication plan should:

- Define and publish terms consistently using industry standard nomenclature across all data gathering tools and processes, reports and recommendations
- Differentiate recommendations, financial models, service definitions, gap analysis and future state proposals between the different components of this initiative: physical facilities (data centers), physical servers, and virtual servers. These terms are frequently co-mingled which causes confusion
- To ensure transparency, inform campus stakeholders about the process and methodologies used to assess the current state of data centers and virtualization
- Provide a clear and concise description of proposed desired states for components of this plan and identify gaps in the proposal requiring more work to define and support recommendations
- Address issues of trust and service. Many campus stakeholders trust third-parties and commercial service providers over internal providers. Transparency in service offerings, metrics, business and service models, and governance will help to address these issues. Clarity in these elements should lead to the creation of a memorandum of understanding and reportable metrics driving accountability
- Campus constituent groups have formed around areas such as data center operations and security. These should be encouraged and supported through the CIO's office as peer relationships help raise the level of expertise and service for all
- Governance models with accountable and visible staff in key positions need to be developed and supported

Training

- Consolidation of physical spaces and computing infrastructure will most likely result in staffing shifts due to efficiencies in the common infrastructure. Staff that no longer need to support servers and facilities should be provided training and certification opportunities in areas of value to them and their departments. This could be

in areas such as software and application support (a pressing need in research areas), supporting high performance and high throughput research computing services, innovative applications of technology, or services for education and outreach
- Staff providing data center services will need to continue to maintain technical certifications from vendors consistent with the technologies in use (e.g. Microsoft, Cisco, Oracle)
- Campus user groups and other formal and informal sharing opportunities will continue to support the sharing of information across the organization and the leveraging of peer resources

Internal Marketing
- Benefits of participating in campus solutions for data center aggregation and virtualization should be promoted as a "carrot" for participating
- Passive "carrots" should be publicized. These might include FP&M not creating new, unapproved, cooling, power and space modifications, boutique networking not being provided to unapproved facilities, and servers not being housed in campus wiring closets

Top-down Communication and Enforcement Needs
- As campus units do not pay for space, power, or cooling for data centers, there is little incentive not to create local facilities
    - Policy is needed to define when a new data center is created, when an existing data center is modified or enhanced, and when a data center should be decommissioned. FP&M staff should not engage in room modifications without this approval. New construction should have a campus review and approval process in place before a data center facility is created (e.g. Division of State Facilities for new construction)
    - "Virtualization first" policies are largely in place across many campus IT service units and should continue to be encouraged by the CIO governance structure and others. There are many systems that are not candidates for virtualization. These can be documented as part of managing the overall campus IT portfolio
    - Data center policies should be consistent across all campus facilities. This can be defined through the data center governance and advisory structure and enforced through the CIO's office. Access controls, auditing, space utilization and planning, networking, facilities management, inventory/asset management, and more will need to be defined. These policies will help to ensure that different organizational units sharing common facilities can work together to efficiently utilize the space and not create risk to the facility, data, network security, devices or other staff

Dependencies or Constraints

The following dependencies and/or constraints were identified as key considerations:
- Incentivizing the right behaviors such as increasing virtualization where possible, utilizing common services/facilities, and utilizing staff expertise to help the broader campus resource. UW-Madison will need to make the cost as reasonable as possible. Common good models, a subsidized core infrastructure, and central funding for key staff positions to avoid charge-back models will be needed
- UW-Madison will need executive buy-in to create a central authority for data center governance
- Policies and procedures will need to have broad support and executive champions to ensure campus-wide adoption
- CIO support is necessary to operationally sustain the governance and service models developed out of this proposal
- While three implementation teams are framed for initial start-up work, the governance team and the facilities team will have on-going policy, process, operations and maintenance responsibilities. The services team will be project-oriented throughout the implementation phase, with on-going work folding into the facilities team. Staff and financial resources should be provided to these teams with this model in mind

- Financial constraints on supporting the relocation of systems, any facilities modifications, and provisioning core service virtual resources have not been determined
- Technical constraints will limit 100% adoption of any solution. These might include vendor-specific configurations, security limitations, and research service or facility-specific issues
- Creating scalable centralized solutions will require the standardization of components and investment in scalable and supportable facilities, hardware, and software infrastructure as an aggregation target
- Key stakeholders will be dependent upon FP&M and Campus Network Services to make aggregation and virtualization successful from a service delivery and performance perspective. Data center planning needs to be closely tied to facilities and network strategic planning
- Staff resources are limited and already fully utilized. Timelines on project components will be dependent upon appropriate staff resources being available in a timely manner, which would mean reprioritizing existing projects and may require staff to change organizational units

Assumptions

In order to prepare this document, certain assumptions have been made:

- Real cost savings as defined in this phase of the AE effort can be realized through the methodologies reflected here (consolidation of servers into a limited number of data centers, virtualization of servers, and by implication, more efficient use of staff resources to provide and support these services)
- Extant staff resources can be combined to provide the central service
- The cost of implementing this solution does not exceed the cost savings threshold desired by the project. UW-Madison will need to spend money to realize the plan. The assumption is that the amount of support required will be considerably less than the financial savings realized from the overall effort
- The implementation team is able to segregate the risk between tenants so problems with one department's devices does not propose problems for others

Project Risks

- If cost savings are not able to be returned to the participating department, ongoing participation and compliance may be at risk
- The data collected by the AE team may not reflect the true number of physical servers present on campus
- A larger than expected portion of remaining physical servers are unable to be virtualized
- A central governance group and service group may lack the ability and resources to monitor and manage server acquisition and facilities expansion across the entire campus
- The central governance group may not gain the trust to manage administrative, academic and research activities in common facilities

Criteria for Measuring Success

Success can be defined in three areas: (1) maximizing utilization of identified and supported campus data centers by distributed campus IT units, (2) maximizing virtualization of infrastructure, and (3) maximizing the expertise of virtualization and facilities staff at the campus level. Specific success elements include:

- As the success of this effort is tied to work of implementation teams, each team should have specific project and communication plans, timelines, and deliverables
- Trust and service satisfaction are frequent themes as to why stakeholders would not utilize more common services. Trust and service parameters need to be defined by the implementation teams, leading to regular surveys of stakeholders and clients about satisfaction with services
- Although campus IT is already highly virtualized and some servers are not able to be virtualized, the team believes there are significant opportunities to realize infrastructure efficiencies through additional virtualization. The campus server inventory should be completed and then monitored over 5 years. This timeframe will

encompass natural hardware and software product lifecycles when moving from a physical to virtual server is ideal. The inventory should show reduced physical servers and increased virtual servers over time

- The current data center inventory should be completed and the documentation of new facilities, modification of existing facilities, and closure of facilities as departments take advantage of consolidated services should be maintained
- There is very little staff effort and energy utilization data for data center facilities on campus. The information available is broad and not generalizable. The implementation teams should define criteria and annually assess these variables to document changes in resource density for providing common services
- On-boarding campus IT operations into common facilities should require documentation of applications, server infrastructure, and staff support. This documentation will yield improved understanding of our campus IT service portfolio, including servers, applications, data, staff, and the impact of leveraging shared resources

Alternatives Considered

From a facilities perspective, UW-Madison has a number of options as it considers the most efficient and effective use of space and resources. The team considered the following scenarios for future-state facilities:

Build a Large UW-Madison Data Center to Supplement and/or Replace the Numerous Existing Data Centers

- Real estate to build such a facility on campus is scarce and valuable. Alternatively, building such a facility off campus, as in Verona or Fitchburg, forgoes the advantage of low-cost cooling from co-gen plants on campus. The projected savings of server consolidation and co-location, as calculated by the accompanying financial model, could not pay for such a building in a 5-yr period. The need for a new data center could possibly disappear even as it is completed if current projections of cloud facilities were to materialize; alternative uses of space built as a data center are few

Outsource Servers, Services and Storage to the 'Cloud'

- Clouds are designed to store data redundantly at multiple data centers, with automatic failover from one to another, thereby preventing a single point of failure in the system as with a single data center. Cloud data centers can be located wherever is most economical and are connected through fast fiber. Third-party commercial sources already offer servers, web hosting and bulk storage at competitive or near-competitive prices. There is institutional reluctance to house research data in cloud facilities that may extend overseas due to legal concerns about export regulations associated with federal grants. File serving and storage in the cloud is currently being negotiated with box.net through Internet2, which is aggressively pursuing additional cloud solutions that are designed to server universities with needs and constraints such as ours, and utilize university consortia-owned fiber to connect. Particularly high service levels, such as PCI for credit card information, could be outsourced to commercial data centers set up for such security requirements rather than trying to retrofit existing facilities or bringing entire facilities to that level when the bulk of their data does not require that security level

Upgraded Campus Data Centers

- Some of the identified 97 server facilities on campus could be candidates for upgrades, increased server capacity and improved utilization by enhancing their power and cooling supply, efficiency, and space management. Dedicated campus-level management could create a 'private cloud' from the aggregated campus infrastructure, where redundancy of data storage could avoid a single point of failure. However, without examining each existing data center in detail it is difficult to know whether the retrofitting of particular facilities is worthwhile. Comparisons must be made not only based on the expense but also in comparison to the alternatives: building new buildings or jumping into the cloud. If properly equipped and managed, the remaining campus data centers could be networked to provide a 'private cloud'

Supporting Documents

The following documents are attached to this business case in support of the proposed solution and related analysis:

- **Appendix 1**: Data Center Administrator Survey – Facilities Database
- **Appendix 2**: Data Center Administrator Survey – Survey Results and Analysis
- **Appendix 3**: Listening Session Notes and Key Takeaways
- **Appendix 4:** Server Scan Summary and Data; Background of Server Scan Process
- **Appendix 5**: Glossary of Common Terms
- **Appendix 6**: Financial Model

Report on Data

A lack of campus-wide data on servers and facilities presented a serious hurdle for the team. While the project timeline did not allow for a comprehensive data collection effort in both these areas, the team initiated a number of widespread data collection processes to support analysis and recommendations, and build the foundation for implementation:

- Scan of campus network segments to provide an objective baseline on current server infrastructure
  - Previous attempts to gather campus-wide data on server numbers, location, types and supporting staff have relied on self-report and have yielded incomplete data sets. In an attempt to provide a more objective baseline, the Office of Campus Information Security (OCIS) was engaged to perform a scan of campus network segments to identify servers
  - The primary tool used to perform this scan was "nmap" (nmap.org). Subnets which were not scanned include Computer Sciences (many servers dedicated to research and not candidates for virtualization or aggregation), ResNET (network segment for students living in residence halls), WiscNET and WiscVPN (not locations of servers within our scope)
  - This list of actual devices active on the network was provided to campus IT leaders for verification and annotation
- Inventory of 97 server room/data center facilities on campus
  - The team engaged with MTAG to help identify data center administrators to invite to a listening session
  - As a follow up to the listening session, the team distributed an online survey to collect server room/data center characteristics. The survey collected data on a variety of attributes
- PUE (power use effectiveness) data was collected by DoIT and FP&M from selected server rooms/data centers to measure how efficiently these facilities use their power

Metadata

Data Item	Data Source Description	Data Structure (Fields)	
Scan of campus network segments	• Data from 25 campus units	• IP's pinged • Hosts Up • Hosts with services • Hosts timeout • SSCC server filter • Total servers	• Physical servers • Virtual servers • Reporting error between department actual server report and SSCC filter of nmap scan
Inventory of campus server rooms/data	• 97 facilities • Data self-reported by server	• Data center location • Associated departments/units	• Physical attributes (e.g. size, weight limit constraints, raised floor,

Data Item	Data Source Description	Data Structure (Fields)
centers	room/data center administrators	• Primary contact • Services provided (e.g. paid hosting, virtual services, backup, web hosting, email) • Primary customer breakdown • Dedicated FTE • Virtualization percentage • Governance (e.g. governance group, external regulations, restrictions on use) • rack system, fire suppression) • Cooling system (e.g. water source, system type, backup) • Power system (e.g. capacity, system type, backup) • Network access • Monitoring • Alarms/alerts • Security arrangements (e.g. card access system, frequency of visits by staff)
PUE data on selected server rooms/data centers	• Collected by FP&M and DoIT	• DoIT • CSSC (computer sciences and statistics complex) – Room B380 • Education – Room L296 Education Building • General Library System – Room 541A Memorial Library • Russell Labs Room B30 • L&S – Room 4411 Social Science Computing Cooperative

Metrics for Future-State Reporting & Analysis

As detailed in the *Criteria for Measuring Success* section, there are several metrics that can be used for future state reporting and analysis. In order for the above-mentioned metrics to be useful the following should be considered for implementation:

- Ongoing surveys of stakeholders and clients regarding satisfaction with services
- Maintenance and monitoring of the campus server inventory, as initiated by the AE team
- Maintenance and monitoring of the campus data center inventory, as initiated by the AE team
- Ongoing assessment of PUE and other energy efficiency data from campus data centers
- Ongoing assessment of level of effort required to support future state infrastructure

Data Accuracy

Data Accuracy / Data Reliability:

- Scan of campus network segments – some departmentally managed or procured firewalls may not have allowed the nmap scanner to perform a complete scan. Thus, the results tend towards underreporting due to the inability of the scanning tool to reach across many firewalls
- Inventory of campus server rooms/data centers – although 97 facilities were identified, this is not a complete dataset (additional outreach required); self-reporting of facility characteristics presents a risk of inaccurate data; some data fields were left blank for a variety of reasons (additional follow-up effort required)

Data Recommendations

The team recommends the following items to improve the suitability, availability, accuracy, and commonality of data for server/data center aggregation:

- Completion of datasets for servers and data center facilities, including both in-person and online outreach

- Implementation team examines the best way to use and maintain these datasets moving forward

Signoffs

Advisory Committee	*Full endorsement received on 5/17*
Steering Committee	*Approved 6/05/12*

AE Initiative Financial Model

Key Assumptions

	Current	2012-2013	2013-2014	2014-2015	2015-2016	2016-2017
Percent of servers virtualized	58%	59%	62%	66%	68%	70%
Percent of total servers in high efficiency data centers	15%	22%	32%	43%	59%	71%

Anticipated Solution Costs – Investment Required to Virtualize and Co-locate

	2012-2013	2013-2014	2014-2015	2015-2016	2016-2017
Cost of servers	$50,497	$144,102	$191,106	$86,242	$148,173
VMware license costs	$25,249	$72,051	$95,553	$43,121	$74,087
VM maintenance costs	$7,575	$29,190	$57,856	$70,792	$93,018
Additional storage costs	$26,511	$75,653	$100,331	$45,277	$77,791
Moving costs – co-location	$55,696	$107,104	$126,839	$137,918	$111,090
Moving costs – virtualization	$31,561	$90,064	$119,441	$53,901	$92,608
TOTAL COSTS	**$197,088**	**$518,163**	**$691,126**	**$437,251**	**$596,767**

Anticipated Savings & Cost Avoidance

	2012-2013	2013-2014	2014-2015	2015-2016	2016-2017
Server replacement costs avoided	$103,926	$407,346	$826,466	$1,059,495	$1,435,741
Annual backup and software costs avoided	$47,527	$186,287	$377,957	$484,525	$656,589
Utility costs avoided	$71,402	$267,865	$471,588	$743,824	$994,086
Labor costs avoided	$28,229	$112,852	$233,788	$305,889	$422,925
TOTAL COST AVOIDANCE	**$251,084**	**$974,350**	**$1,909,799**	**$2,593,733**	**$3,509,341**

Five-Year Financial Projection

	2012-2013	2013-2014	2014-2015	2015-2016	2016-2017
Variable costs	$197,088	$518,163	$691,126	$437,251	$596,767
Cost avoidance	$251,084	$974,350	$1,909,799	$2,593,733	$3,509,341
GROSS SAVINGS	**$53,995**	**$456,187**	**$1,218,673**	**$2,156,482**	**$2,912,574**
Up-front costs*	TBD	TBD	TBD	TBD	TBD

*Initial capital investment (if necessary) to be determined by implementation team

Advisory Committee Meeting: May 17, 2012 10:00am-11:30am

Discussed Q&A:

Question *– Please explain the TBD up-front costs to upgrade facilities. How does the team propose to quantify upfront costs?*

Answer – Deciding whether to invest in all or some of the options including in-situ upgrades, retrofitting, new construction, and public-private partnerships requires a degree of data collection and analysis that extends beyond the AE team's work. The AE team proposes the following steps be taken by the implementation team to estimate the capital investment to support campus' future-state server infrastructure:

1. Quantify space required – detailed assessment of demand for space considering the following variables:
    a. Degree of virtualization
    b. Server growth rate
    c. Movement to cloud-based services
    d. Technological advances allowing for improved density
    e. Demand for computing and storage
2. Assess space available:
    a. Determine potential to scale/upgrade specific facilities
    b. Identify space available for new facilities
3. Evaluate cost/benefit of alternative options:
    a. Review entire option set for on-campus/off-campus solutions
    b. Assess implementation approaches (e.g. pilot, staged, opt-in)
    c. Conduct actuarial analysis to understand costs and benefits
    d. Understand existing capital plans for facilities upgrades to determine capital avoidance
    e. Determine operating expenditures to run/scale targeted facilities

Question *– How will potential customers become aware of the campus shared data center service provider?*

Answer – The team has identified potential early adopters that a) want nothing to do with the physical server boxes and/or b) want assistance with virtualization or co-location. Broader campus engagement will likely be facilitated by word of mouth via effective engagement with this early adopter customer group. An effective communication plan outlining features, services, and benefits will also be crucial to widespread adoption.