Administrative Excellence
UNIVERSITY OF WISCONSIN–MADISON
Shaping our Future

Project Team:	**Administrative Excellence**
	Data Center Aggregation Work Team Project Agenda
Date / Time:	Tuesday, April 24, 2012,
Location:	Room 362 Memorial Library
	10 am – 12 pm
Attendees:	Ed Van Gemert, Steve Krogull, Phil Barak, Kevin Cherek, Nancy McDermott, Melissa Amos-Landgraf, Chris Slatter,

----- AGENDA ----

TOPICS	WHO	TIME
Review notes from April 10, 2012	All	5 min.
Identify data center service layers/needs **Match service layers with current service level offerings from identified data centers. **Identify what is needed to upgrade selected facilities to cover service layer needs **Identify derivative work.	All	30 min.
Facilities: What do we need to know about these selected facilities? Capacity? Renovation needs? How does building something fit into our planning? How much is derivative work?	All	30 min.
Break		5 min.
Review/discuss draft of the proposed solution description section of the business case.	Kevin C.	30 min.
Server inventory: review current available data	Steve K.	15 min.
Wrap up, next steps	All	5 min.

AE Business Case Status – Data Center Aggregation		
Business Case Element	**Assigned Team Member**	**Status**
Initiative Sponsorship and Ownership **Project Name:** AE IT Data Center Aggregation **Project Summary:** Analyze and identify potential cost savings and efficiencies including: utility costs, labor efficiencies and effective best data center practices, space re-utilization costs, and improved risk management and security measures through greater virtualization, consolidation and co-location of servers and data centers across campus. **Business Unit(s):** Vice Chancellor for Administration and Budget—Administrative Excellence. **Business Process Owner(s)** TBD **Primary Cost Estimate:** To be supported by a separate financial model document. **Proposed Go-Live Date:** TBD—based on proposed solution milestones and timing.	Ed V.G.	Completed 3/6/12
Business Need or Opportunity A thorough current state campus server inventory, data collection including network scans, power utilization effectiveness (PUE), focus groups, and data center survey analysis is needed in order to identify and recommend future campus server/data center needs and services. The opportunity exists today to encourage best data center practices using incentives that will result in significant cost savings for central campus, individual schools and colleges and administrative units through the reorganization and consolidation of server placements and solutions including increased	Ed V.G.	Completed 3/6/12

virtualization, co-location and outsourcing		
Alternatives Considered ***Scenarios for UW-Madison Data Center Aggregation for Admin Excel Team*** *Disclaimer: None of the following scenarios are recommendations, either personal or team, and certainly not plans. Some may be regarded as straw-man arguments constructed simply to be thorough in deliberations.* **1. OutSource Servers and Storage (aka, the U.S. Cloud)** Case: Building and maintaining server and storage facilities in Madison, where real estate is expensive and with no local energy sources except for those imported from other regions, seems counter-intuitive. As an alternative, let other facilities attached to the fast fiber connections bid for our business. Offshoring would involve unnecessary entanglements with export restrictions for intellectual property but outsourcing to, say, Montana (coal), Tennessee (hydroelectric) or elsewhere within the US, might suit. Migration/Transition: The servers in the least secure facilities would be the first to migrate, perhaps at little or no cost to users so as to minimize resistance to change. Similarly, old servers would be migrated instead of replaced to minimize replacement costs. Network scans would verify that new servers and server locations didn't pop up in their place. Ultimately, all servers would be migrated out of Madison during a normal replacement cycle, perhaps five years. Only high throughput and high performance computing services would likely remain on campus, if individual business cases or evidence of unsuitability could be produced. Local server administrators would transition from tending the hardware to providing customer services with the applications on the remote hardware. Contra: Investments in existing campus data centers—Dayton St, Genetics, Education, etc.—would be lost, unless power savings offset. Some of that space would not easily revert to offices or labs. **2. Single OffCampus Data Center (aka, the Wisconsin Cloud)** Case: Building and maintaining server and storage facilities on campus, where real estate is limited seems counter-intuitive. As an alternative, utilize the fast fiber connections to west Madison to permit construction of a new server and storage facility from the ground up, perhaps in Fitchburg or Verona. Consider best practices: free cooling or geothermal cooling to reduce energy costs. Migration/Transition: The servers in the least secure facilities would	Phil B./Nancy M.	Completed 2/6/12

be the first to migrate, perhaps at little or no cost to users so as to minimize resistance to change. Similarly, old servers would be migrated instead of replaced. Network scans would verify that new servers and server locations didn't pop up in their place. Ultimately, all servers would be migrated off campus during a normal replacement cycle, perhaps five years. Only high throughput and high computing services would likely remain on campus, if individual business cases or evidence of unsuitability could be produced. Local server administrators would transition from tending the hardware to providing customer services with the application on the hardare.
Contra: Investments in existing campus data centers—Dayton St, Genetics, Education, etc.—would be lost, unless power savings offset. Some of that space would not easily revert to offices or labs. The new facility would be a single point of failure in the event of disaster.

3. Single OnCampus Data Center (Campus Cloud)

Case: None of the existing campus data centers is large enough to meet all campus needs for servers and storage. A single data center on campus would have to be properly sized for space, power and cooling.
Migration/transition: much as 2) above, but migrated on campus to new data center.

Contra: Large localized demand for power and cooling. Adequate space and utilities are unlikely to be found in existing campus building, requiring new building. Such as facility would be a single point of failure in the event of disaster. Investments in existing campus data centers—Dayton St, Genetics, Education, etc.—would be lost, unless power savings offset.

4. Distributed OnCampus Data Centers

Case: A number of data centers and server rooms exist on campus and demand exists for upgrading server rooms into data centers. Servers and storage can be moved from lowest quality facilities into the existing quality data centers. If existing data centers of requisite quality cannot accommodate all of campus servers and storage, then either they will be expanded/upgraded or the best of the server rooms will be upgraded into data centers. Perhaps 10-15 data centers may remain.
Migration/transition: like 3) above. Some of the data centers have lease rates per rack unit or per rack that would be best standardized and/or subsidized to reduce resistance to change.

Contra: Almost certainly the best 5-10 data centers on campus cannot contain all the campus servers and storage without serious rehab, particularly since power and cooling upgrades will be required.

5. Augmented OnCampus Data Centers Case: Once the existing 5-10 best quality data centers on campus are filled with servers and storage co-located from the lowest quality space or most matching the existing data security standards, augmenting capacity with another built-from-the-ground-up data center, either on campus or off, could be used to meet the additional requirements. Migration/transition: like 4) above Contra: Some of the presumed labor savings from having a single site with a small, efficient crew of server administrators may be lost with numerous sites but perhaps inefficiencies could be mitigated by remote access and surveillance of all data centers by a single crew. **Observations:** Savings of labor and power/cooling are achieved in two different but related processes: virtualization of servers and co-location of servers. Virtualization is basically consolidation of virtual machines on a physical host. Absolute limits are set by CPU utilization--operations that are computationally intense cannot be consolidated as densely, if at all, on a physical host as those that are less intense, such as simple web hosting. Virtualization has been underway on campus for some years, with varying degrees of completion of the process. Many server rooms are unoptimized, as when the room is cooled instead of the racks themselves. In such cases, energy requirements for cooling, fans, and uninterrupted power supply may be as much as double the power requirements of the servers themselves. Optimized data centers can have in-rack cooling or hot&cool zones; excess heat, though of low heat differential, can be seasonally utilized for conditioning air intake for additional power efficiency. Proper data centers are also more likely to have formal data security standards that allow them to be assigned a service tier suitable for the functions served. Proper sizing of facilities to be built or rehabbed may be difficult to guess but essential to approximate accurately. For example, co-locating a number of unvirtualized servers into a data center without virtualizing them first could lead to building a larger data center intended to receive them than if they were first virtualized. Conversely, co-locating a number of servers that were already virtualized or unable to be further virtualized because of the extent of cpu utilization could lead to under-building if overestimates were made a priori about the possible level of virtualization achievable with those servers.		

Administrative Excellence
UNIVERSITY OF WISCONSIN–MADISON
Shaping our Future

Proposed Solution Description	**Kevin C. = lead**	
• Scope of Proposed Solution	*Kevin/Rick K.*	
• Proposed Milestones and Timing	*Kevin/Phil B./Dan*	
• Alignment with Strategy	*Kevin/Nancy M.*	
• Customer Readiness	*Kevin/Melissa/Dan*	
Impact	**Melissa A-L = lead**	
• Anticipated Benefits	*Melissa/Phil B./Dan*	
• Stakeholders Impacted	*Melissa/Nancy M.*	
• Impact on Other Initiatives	*Melissa/Steve K./Dan*	
Project Success Factors	**Steve K. = lead**	
• Change Management Plan / Communication Plan	*Steve/Melissa*	
• Dependencies or Constraints	*Steve/Phil B./Dan*	
• Assumptions	*Steve/Rick K.*	
• Project Risks	*Steve/Rick K.*	

• Criteria for Measuring Success	*Steve/Phil B./Dan*	
Supporting Materials	All	
Report on Data	**Chris S. = Lead**	
• Metadata	*Chris S. / TBD*	
• Data Accuracy	*Chris S. / TBD*	
• Data Recommendations	*Chris S. / TBD*	
Financial Model	Phil B./Rick K/ Ed V. G./CS	

File name: DataCenterAggregationagenda042412