

Some Joules Are More Precious Than Others: Managing Renewable Energy in the Datacenter



Christopher Stewart
The Ohio State University
cstewart@cse.ohio-state.edu

Kai Shen
University of Rochester
kshen@cs.rochester.edu

Clean, Green Renewable Energy



Green House Data

- 10,000 sq. ft. in WY
- Wind powered

Similar datacenter plans in
CA, TX, MA, NY & NE

■ *Benefits: Green and Green*

Reduce carbon footprint
compared to Grid (now)



- The Grid's future is uncertain, but a long way from being mostly clean
- Coal still 2X in 2030 [EIA, 2009]

Profitable, **Green** Renewable Energy



AIso.net

- Solar-powered datacenter in CA
- Competitive monthly hosting \$9.95-\$50
- Profitable

■ *Benefits: Green and Green*

Equipment produces a lot of energy; low maintenance
° human costs



- Energy storage not as important as in Grid
- Incremental growth
 - [Google's Solar Panel Project]
 - Buy equip. or pay as a service
- Amortize cost over equip. lifetime

The Intermittency Problem

Applications in the datacenter must be available 24x7

But wind and solar energy are *intermittent*

- Datacenters powered by renewable energy need backups
 - Primary options: Grid, generator, battery
 - Alternatives are either dirty and/or costly

- *Renewables are precious!*
 - renewable = joule converted from solar/wind
 - The preferred energy source
 - Available only sometimes and costly to store

Position Statement

1. Datacenter design and management should strive to increase the use of renewables

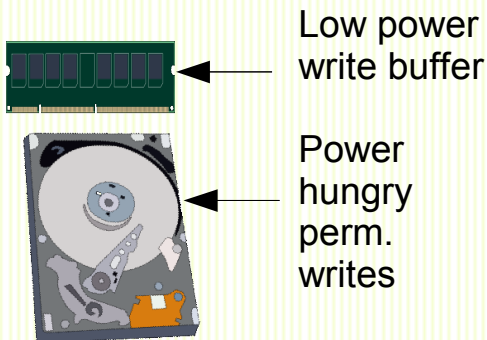
Opportunities

- Capacity planning
 - Compute power should fluctuate with intermittent outages—i.e., turn machines off
- Load balancing
 - Route requests to datacenters with unused renewables
 - [K. Le-HotPower-2009] [A. Hopper-2008]
 - Migrate services to datacenters with renewables

2. Need mechanisms for renewable-aware management

Intermittent Renewables are More Precious

- Every Joule is Precious[Vahdat, 2000], right?
- Energy efficiency alone may not be enough
 - Consider the architecture of a storage device



Policy w/ Energy Efficient Management:

1. Put the disk to sleep
2. Fill the low power write buffer
3. Dump write buffer all at once, in a bursty fashion

What if renewables are available during the sleep?

A Renewable-Aware Policy

1. Put the disk to sleep
2. Fill the low power write buffer
3. Dump write buffer whenever renewable energy is available

If write buffer fills and no renewables are available, do we fall back to energy efficiency?

Outline

- Studying the Impact of Intermittency
 - Model of a wind-powered datacenter
 - Effect on wind energy production
 - Use of renewable energy in the datacenter

- Managing Renewables at the Request Level

- Conclusion

Intermittency

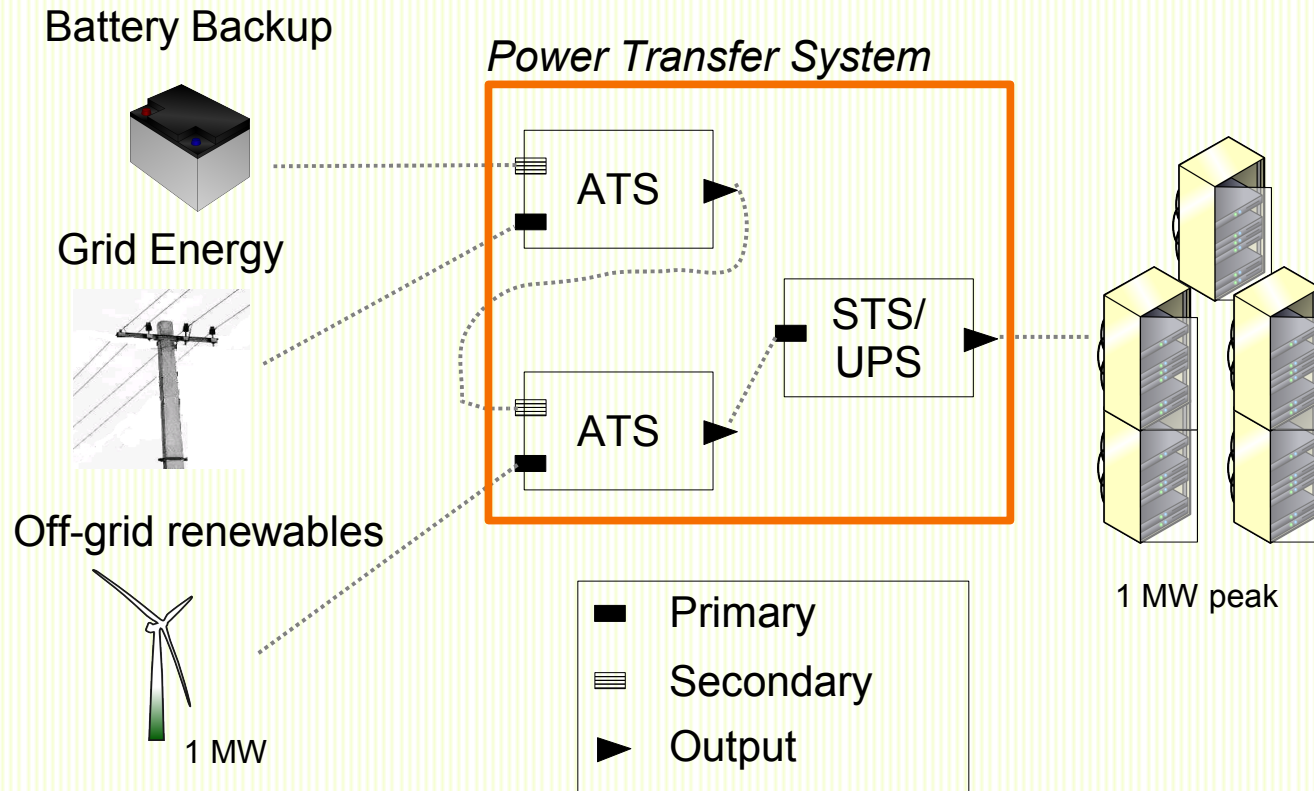
1. Datacenter Model
2. Wind Intermittency
3. Renewables in the DC

Approach

- Understand the flow of power through a datacenter
 - Power sources, e.g., the Grid or a renewable energy
 - Power-transfer devices, e.g., ATS, PDU
 - Power consuming equipment
- Identify parameters that affect the use of intermittent power from renewable sources
- Assess the impact of these parameters

Intermittency

1. Datacenter Model
2. Wind Intermittency
3. Renewables in the DC



Power transfer system:

Input: K High voltage, turbulent AC streams

Output: N high-quality power streams to N compute devices

Intermittency

1. Datacenter Model

2. Wind Intermittency

3. Renewables in the DC

■ Automatic transfer switch (ATS)

- Input 2 power sources, Outputs 1 power source
- Prevents dangerous “backfeed” between sources
 - Only one source is active at a time
- *Here's how it works*



1. Monitors power from the primary source

2. When power from primary dips below K , ATS switches (with no detectable delay) to secondary

3. When primary exhibits K power for K' seconds, ATS switches back to primary

- Primary
- ▨ Secondary
- ▶ Output

Intermittency

1. Datacenter Model

2. Wind Intermittency

3. Renewables in the DC

Key parameters related to the ATS

■ Power-transfer threshold

- Power level where ATS switches to the secondary source
- For dependability, threshold equals peak consumption
- For high use of renewables, set the threshold low

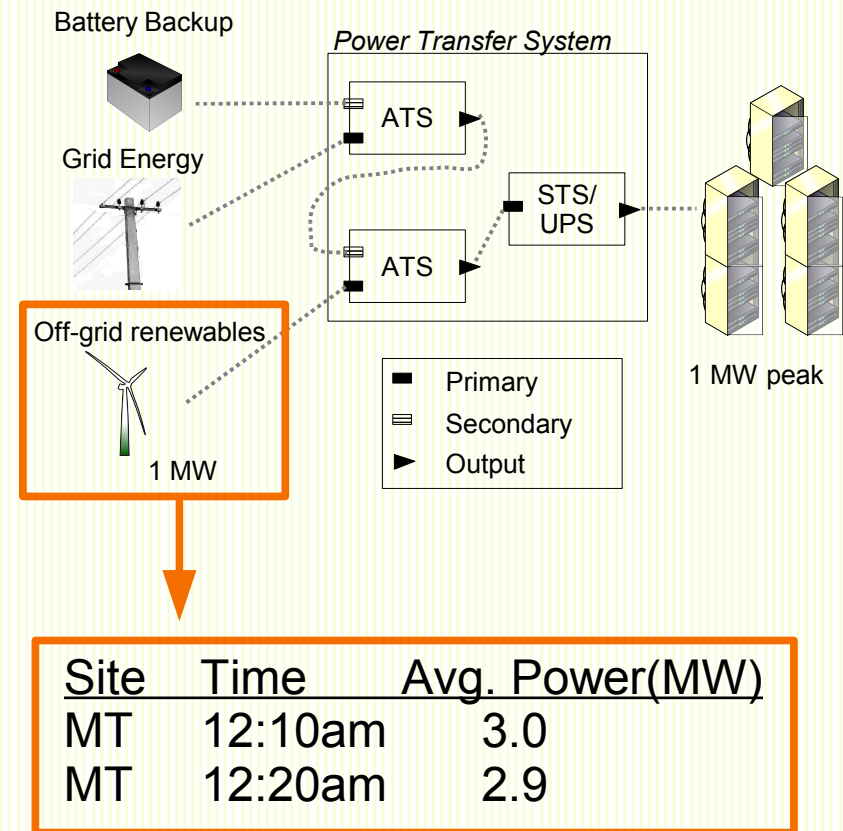
■ Average power consumption of the datacenter

- When renewables are available
- In relation to the power-transfer threshold
- Research challenges: Usefully increase power density (i.e., do more work) without degrading performance

Intermittency

1. Datacenter Model
2. Wind Intermittency
3. Renewables in the DC

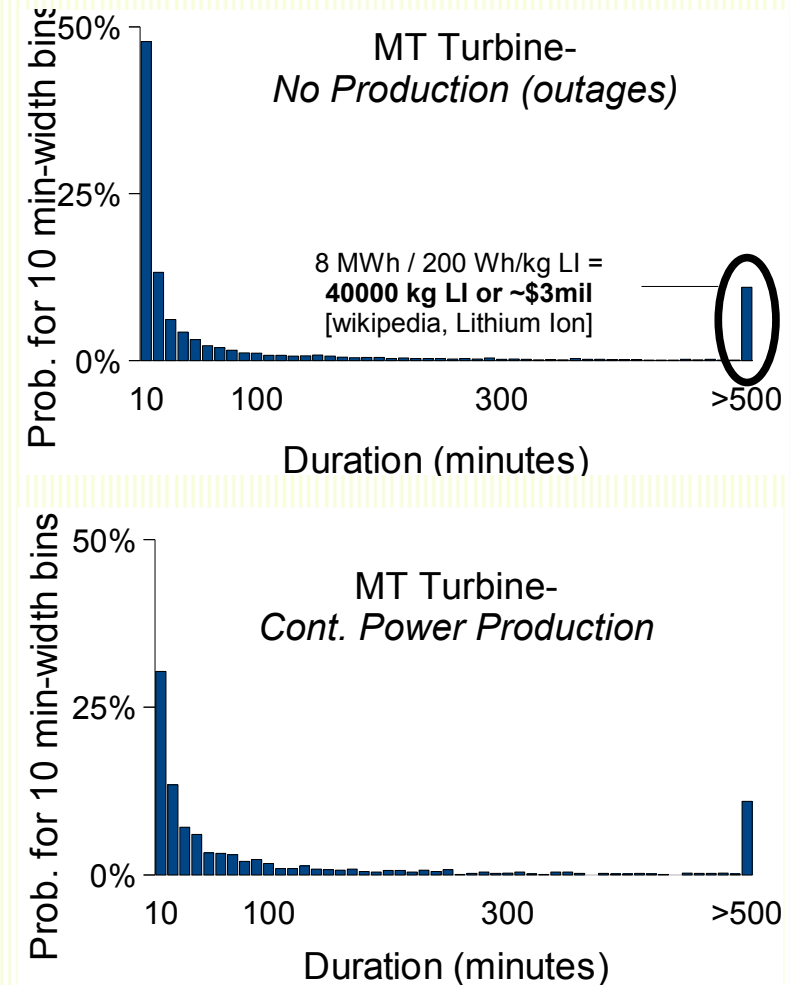
- Goal: Understand impact of
 - 1. power-transfer threshold
 - 2. power consumption
- Approach: Simulate the behavior of the ATS under different datacenter consumption scenarios
- Challenge: Intermittency Renewable Sources
 - We used traces from National Renewable Energy Lab
 - 10 min. snapshots of power production at turbines in MT and CA
 - Normalized to a 1 MW turbine



Intermittency

1. Datacenter Model
2. Wind Intermittency
3. Renewables in the DC

- *How precious are renewables?*
- Plot intermittent outages and continuous power
- **Battery backup too costly**
- **But enough continuous supply for renewable-aware management**
 - Some power production 42% of the time for MT turbine
 - Peak power 25% of the time
 - *Peak renewables are 40% percent more precious*



Intermittency

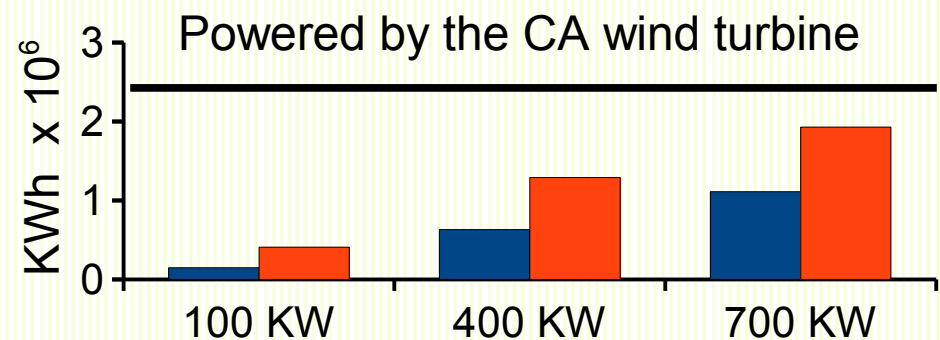
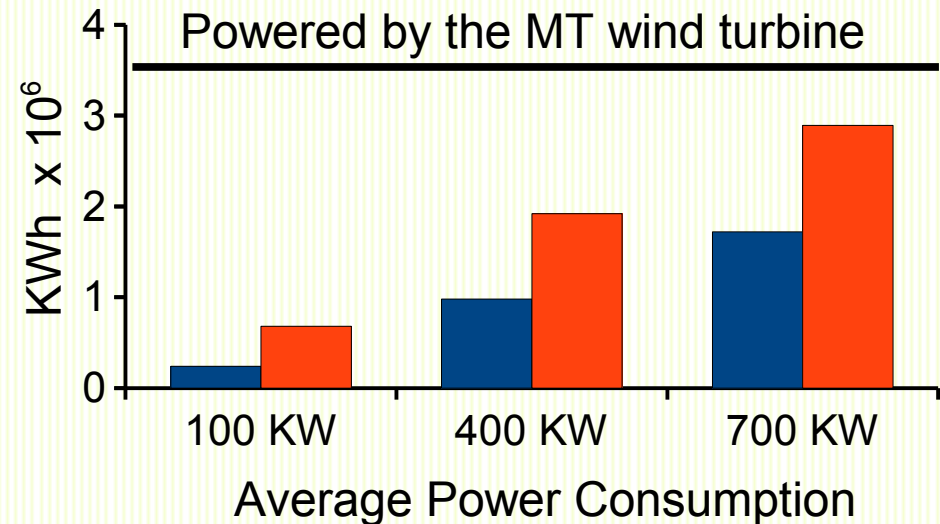
1. Datacenter Model
2. Wind Intermittency
3. Renewables in the DC

- NREL data provided input to our datacenter model
- **Metric: Renewable energy used**
 - MT produced 3.5M KW-hours of renewables and CA produced 2.3M KW-hours (black line)

■ Power-transfer threshold = peak power

■ Power-transfer threshold = zero

- **1 MW Power-transfer threshold caused up to 65% drop**
- **Zero power-transfer threshold complicates the benefits of increased consumption**



Intermittency

1. Datacenter Model
2. Wind Intermittency
3. Renewables in the DC

- Metric: Cost per KW-hour
- Average price for commercial electricity \$0.10 KW-hour
- \$2.4M to erect a wind turbine that is connected (directly) to a datacenter [European Wind Energy Assc.]
 - \$1.6M installation
 - 2% annual maintenance fees
 - Lifetime of turbine: 20 years
- Datacenter at CA or MT could use 24M KWh
 - Either high power consumption or zero threshold
 - Wind-powered datacenter in MT: *\$0.04 KWh*

Outline

- Studying the Impact of Intermittency

- Managing Renewables at the Request Level
 - Request-Level event profiling [shen-asplos-2008]
 - Event-driven power modeling [bellosa-ew-2000]
 - Energy and Power Characteristics of Requests
 - Load balancing scenario

- Conclusion

Managing Renewables

1. Request Characteristics
2. Scenario
3. Future Work

Prev. section: MT wind-powered Datacenter with high average power consumption (when renewables are available) can potentially reduce costs to \$0.04 Kwh

- Idea: Use request-level management to increase power consumption when only renewables are available
 - Route more requests to datacenters with unused renewables
 - Delay long-running Map-Reduce requests
- Need mechanism to predict power/energy consumption of a request

Managing Renewables

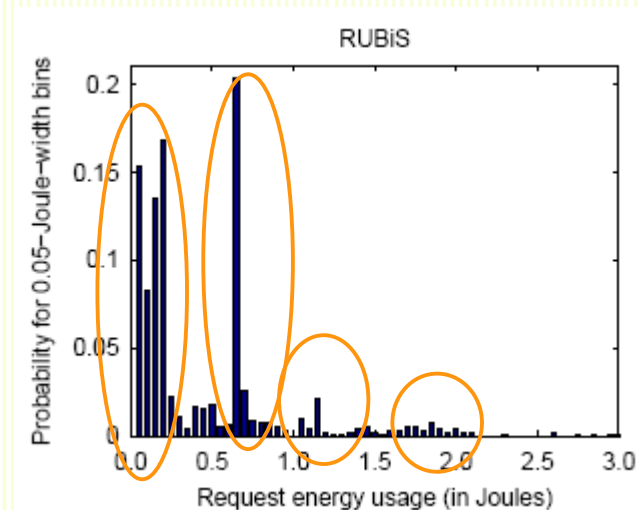
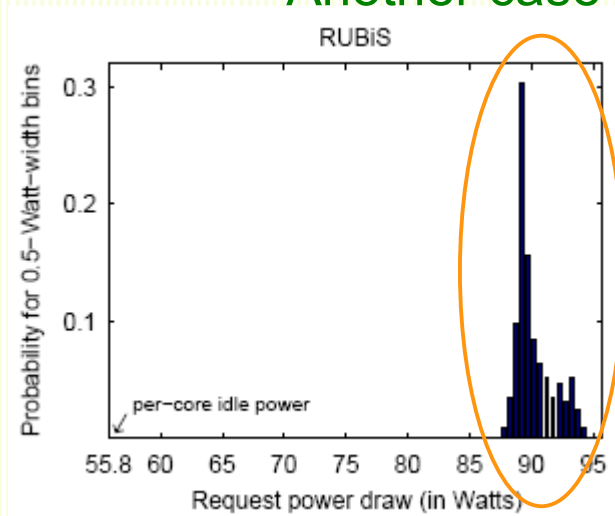
1. Request Characteristics

2. Scenario

3. Future Work

- Request workloads executed in isolation
- WattUp power meter measures watts and joules
- Processor was not adjusted during tests
 - Disk accounted for as static consumption

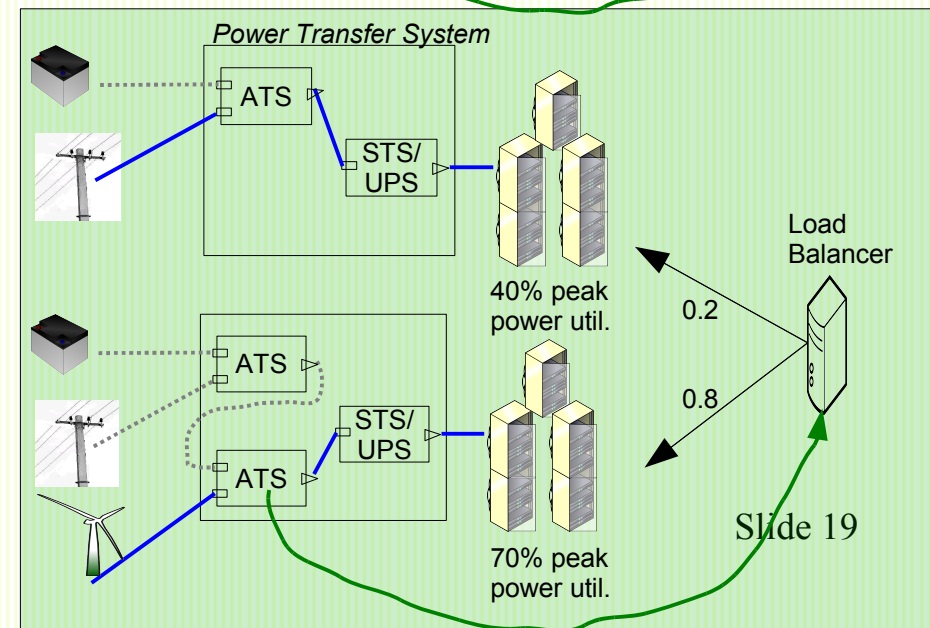
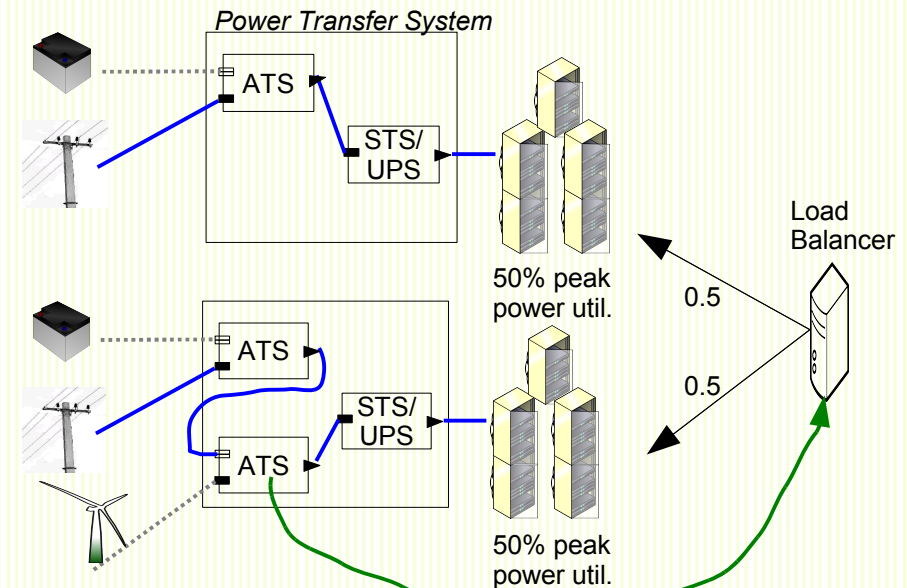
Another case for energy proportional servers



Managing Renewables

1. Request Characteristics
2. Scenario
3. Future Work

- Consider 2 futuristic datacenters
 - Server power is energy proportional
 - 1 datacenter uses renewable energy
- Power monitoring system notifies when precious renewables are available
- Load balancer increases the renewable-powered datacenter's workload
 - Uses request power models as guidance
- Challenges
 - DCs connected by hard-to-change DNS
 - Need optimization procedures [K. Le-HotPower-2009]



Managing Renewables

1. Request Characteristics
2. Scenario
3. Future Work

Our ongoing research is on the design, management, and mechanisms that will increase the use of *precious renewables in the datacenter*

1. Study trade offs in the configuration of the power-transfer system, i.e. ATSS, PDUs, and UPSs
 - Amount of precious renewables transferred to compute equip.
 - Brownouts
 - Practicality—space, cost, and increment-ability
2. Study load balancing policies with renewable-aware mechanisms
 - e.g., per-request vs. per-session round robin, feedback based
 - Redirection speed for intra- and inter-datacenter workload