

CICADA

PREDICTIVE GUARANTEES
FOR CLOUD NETWORKS

Katrina LaCurts

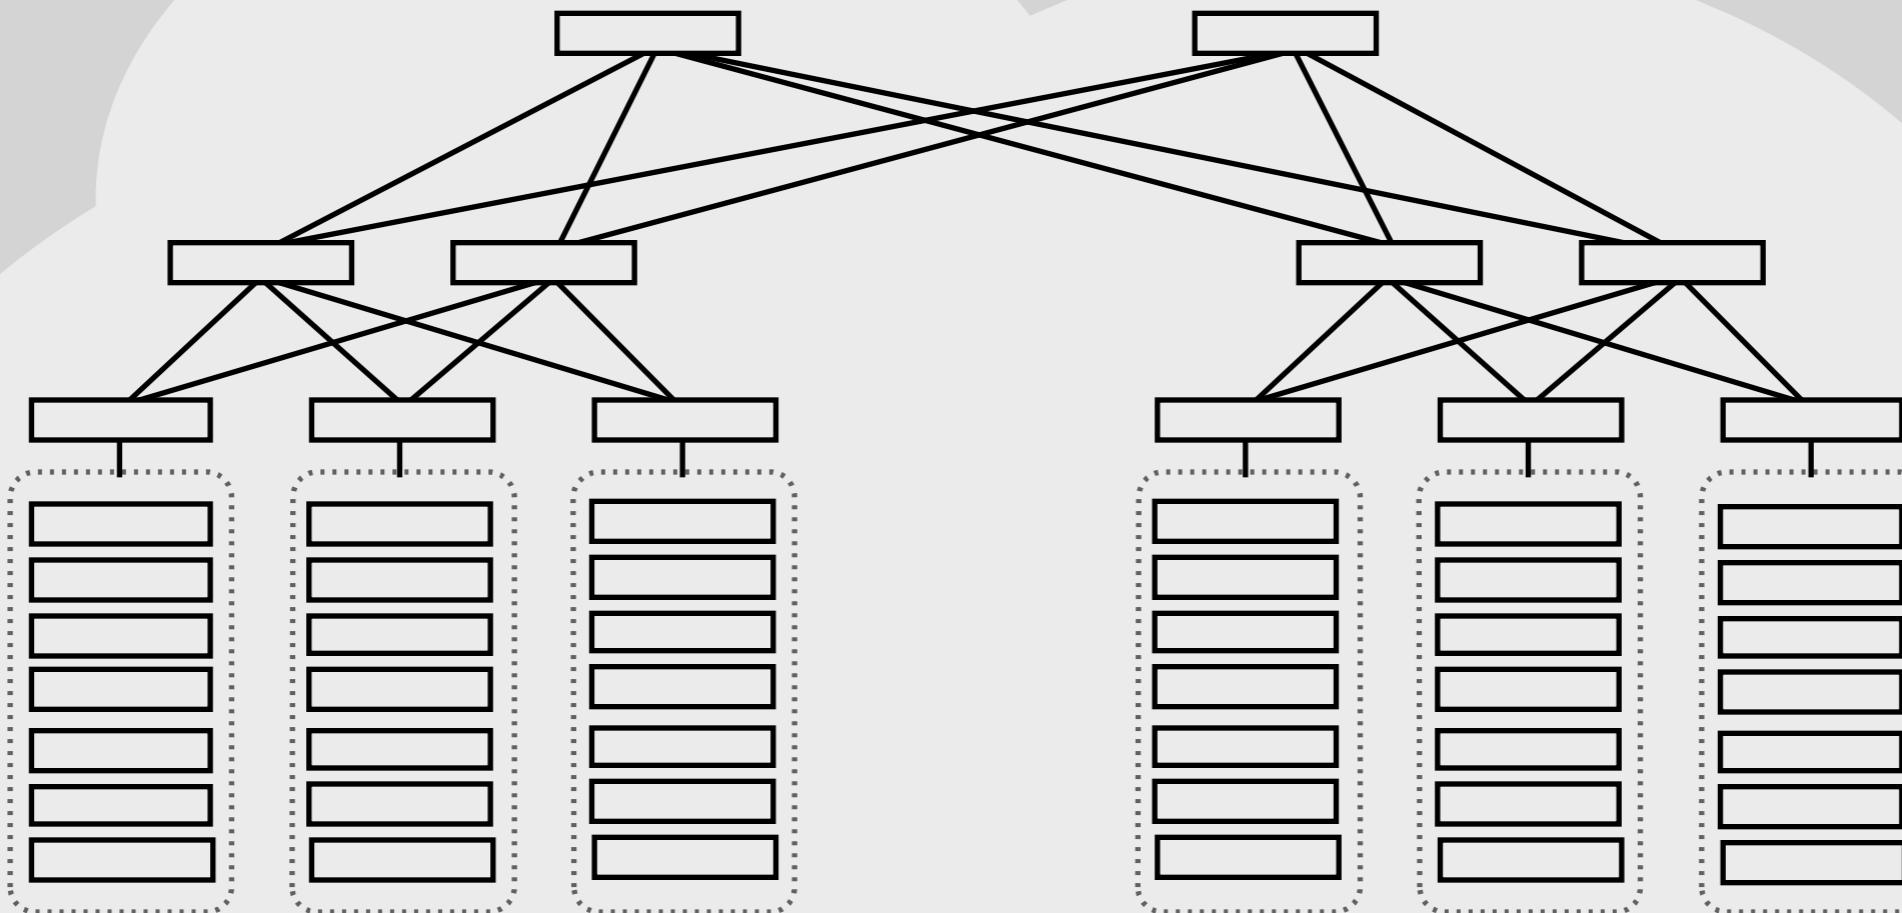
Jeff Mogul

Hari Balakrishnan

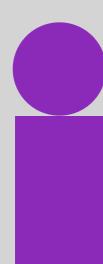
Yoshio Turner

MIT CSAIL, Google Inc., HP Labs

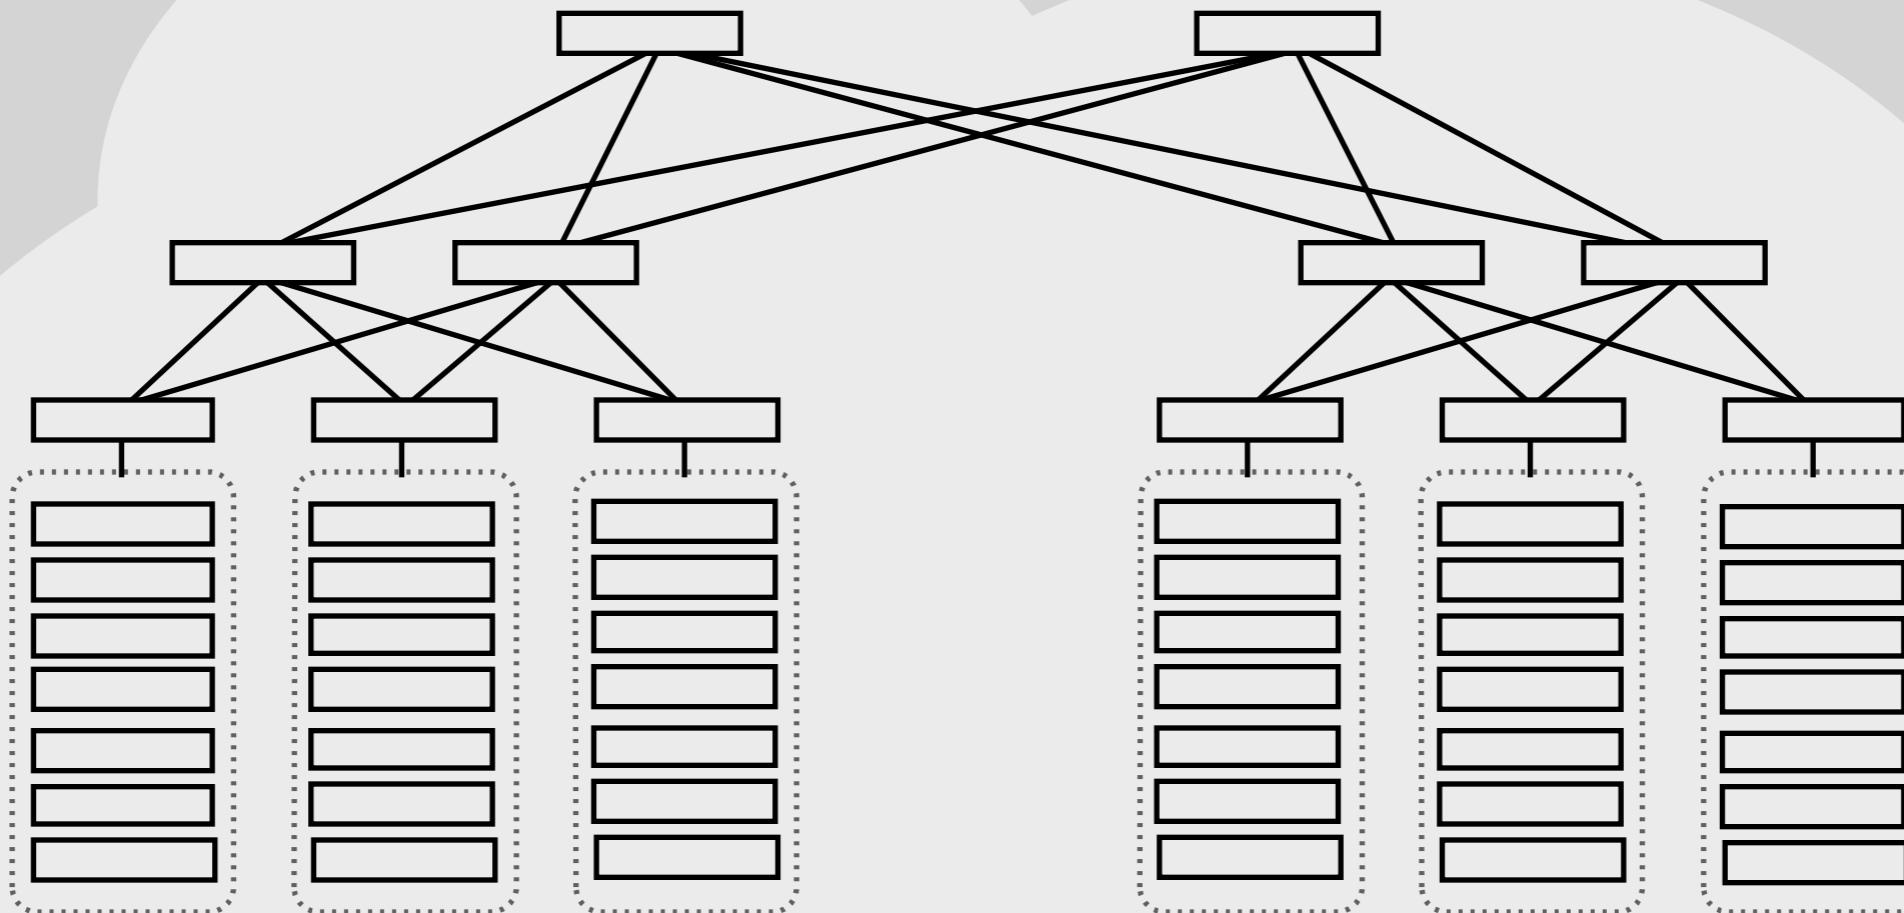
BANDWIDTH GUARANTEES



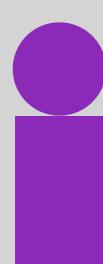
BANDWIDTH GUARANTEES



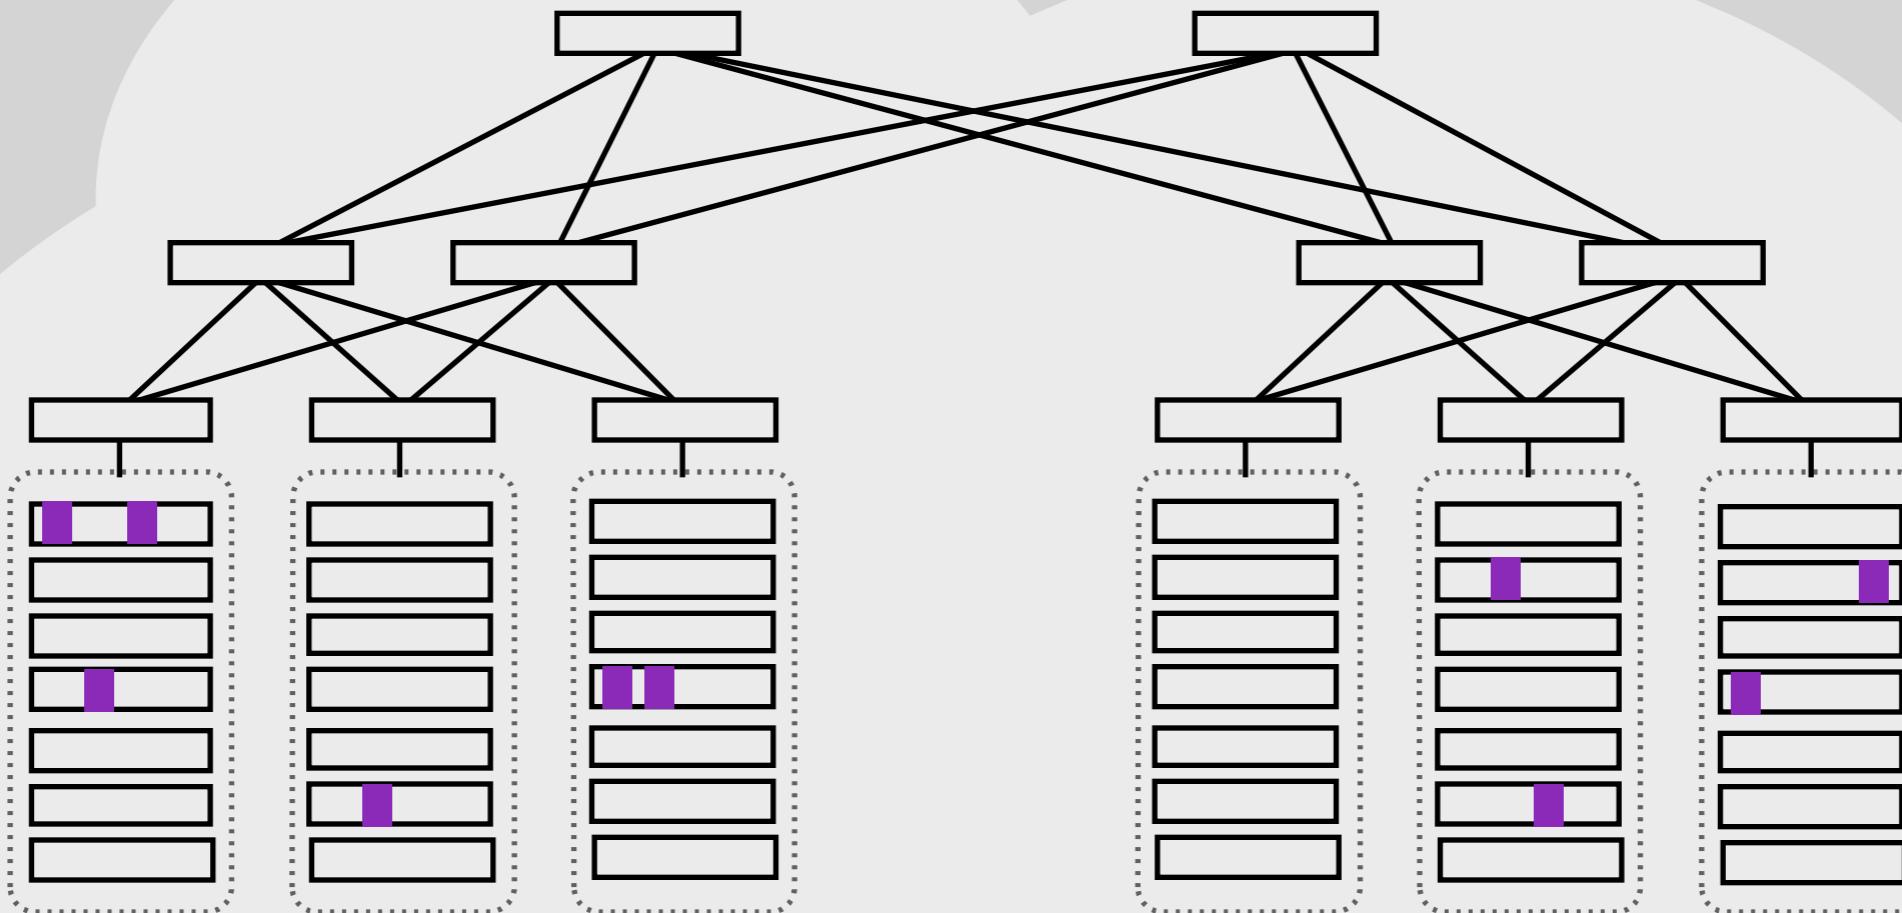
10 machines



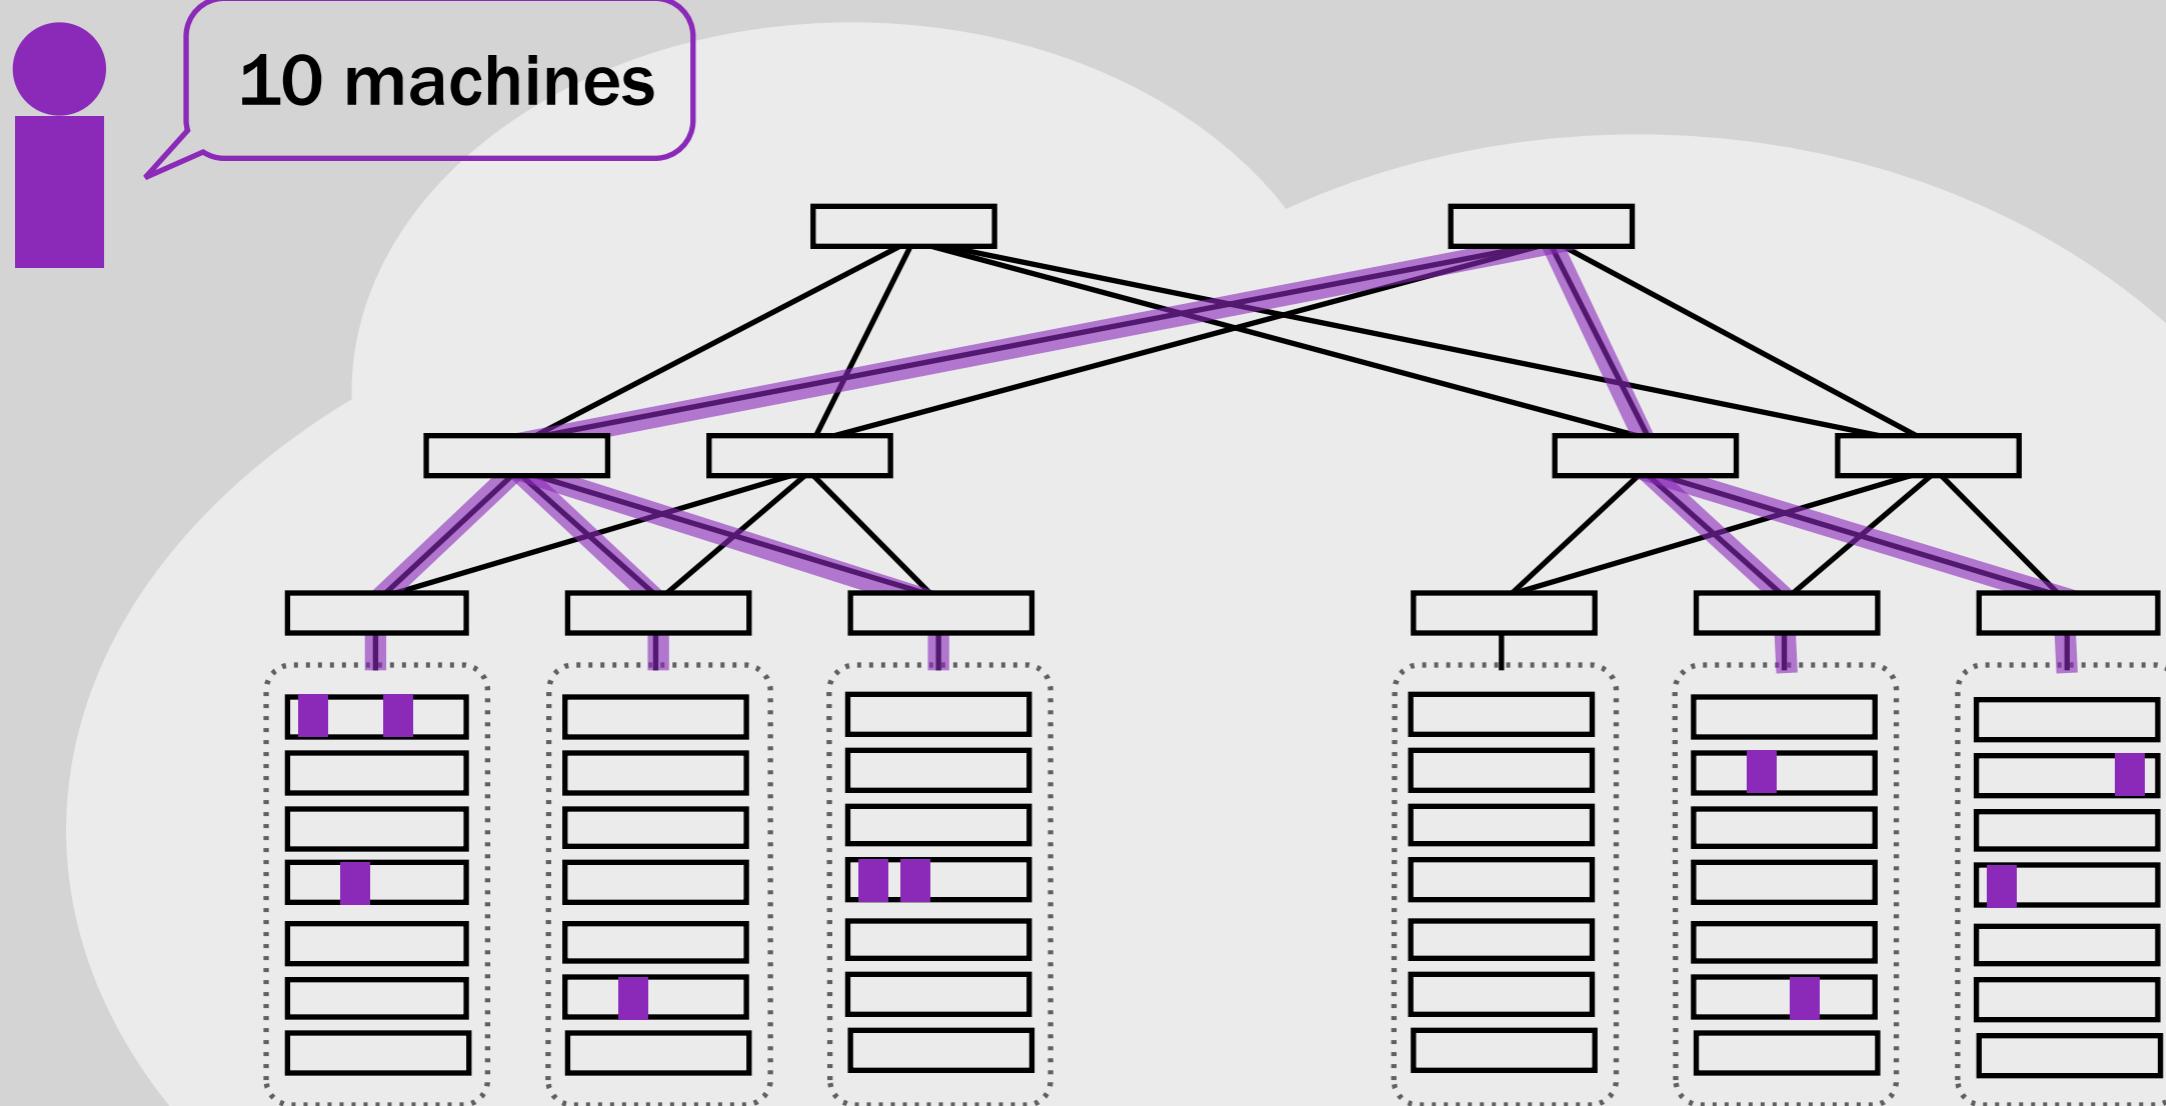
BANDWIDTH GUARANTEES



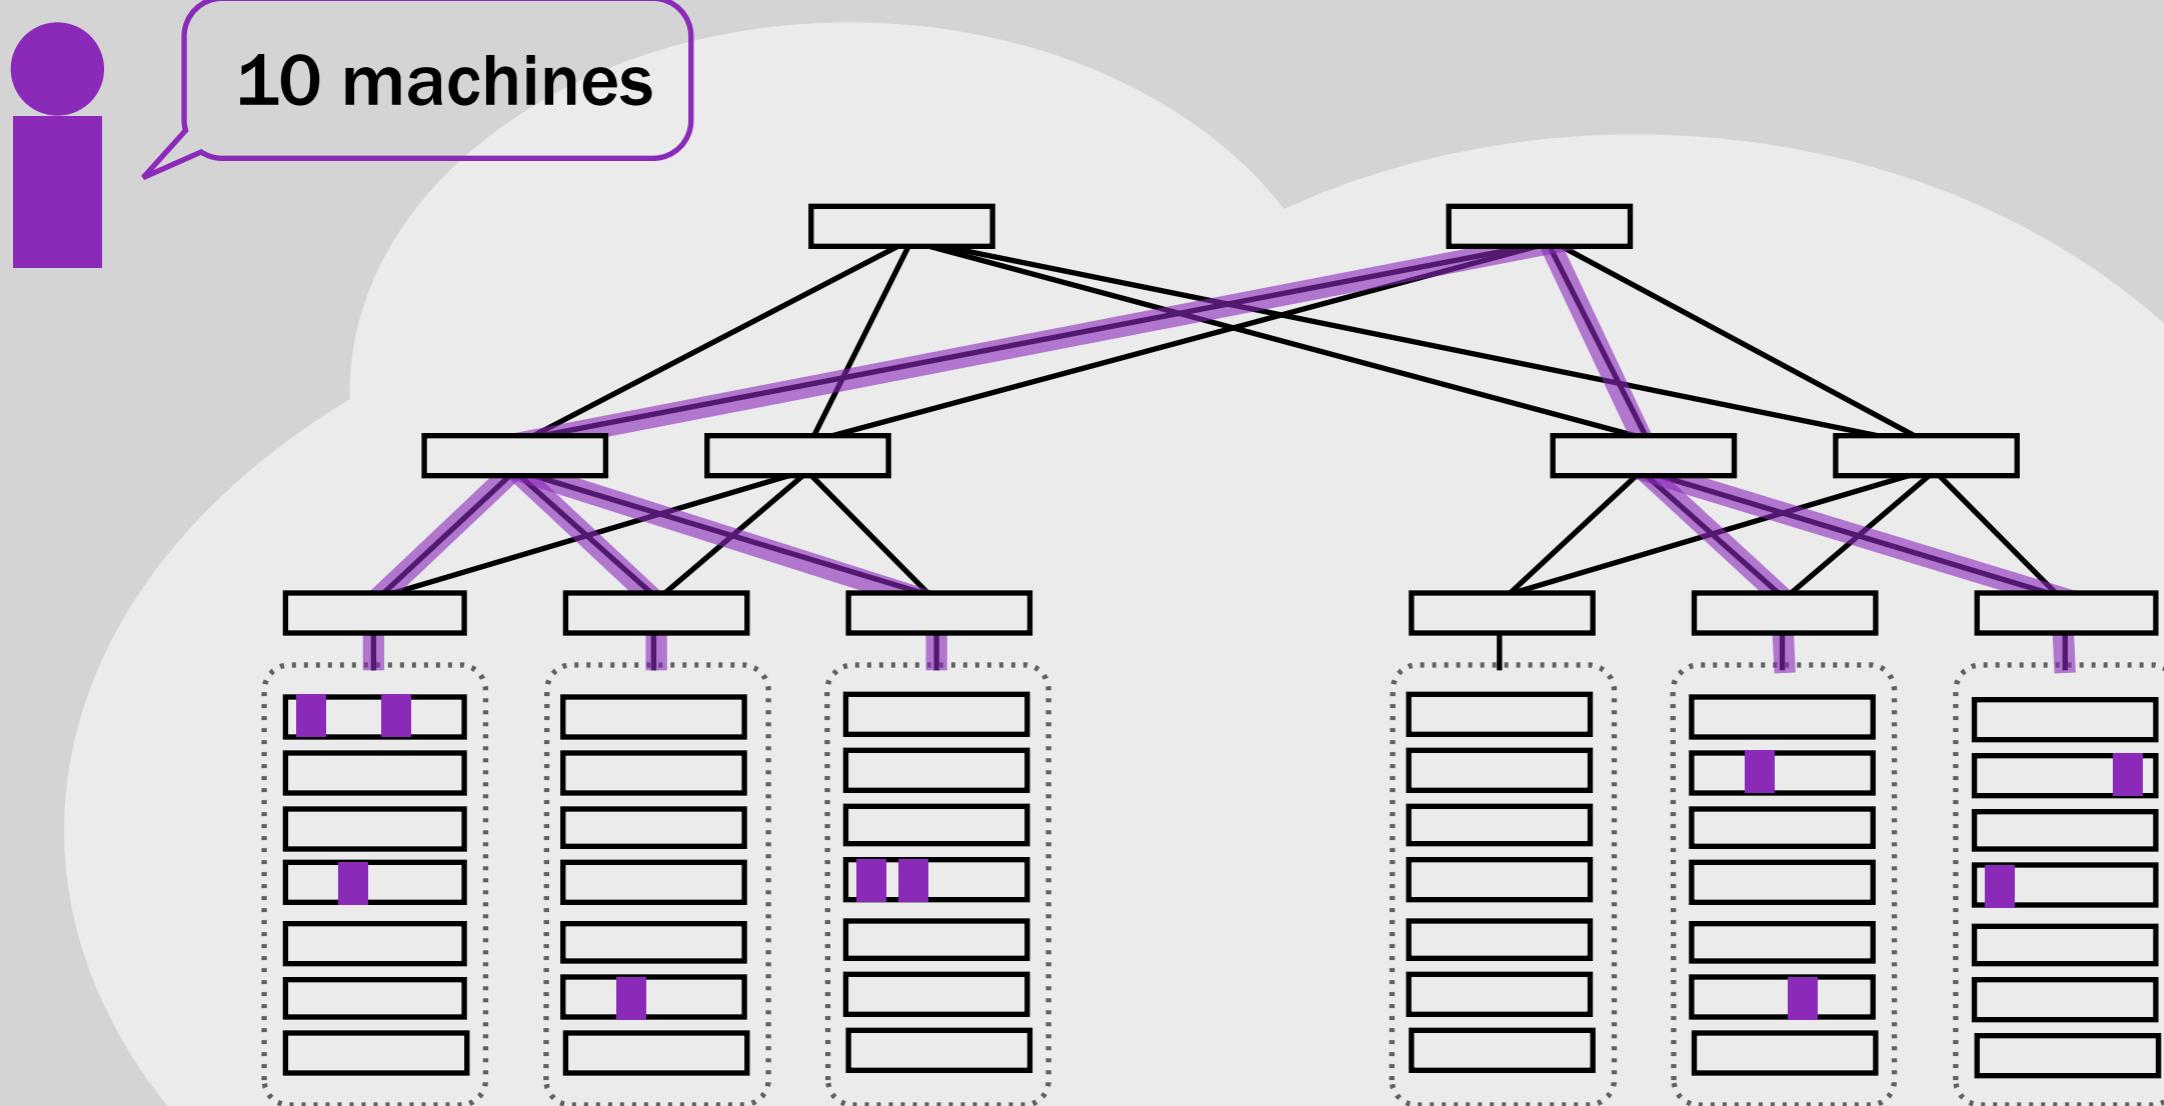
10 machines



BANDWIDTH GUARANTEES



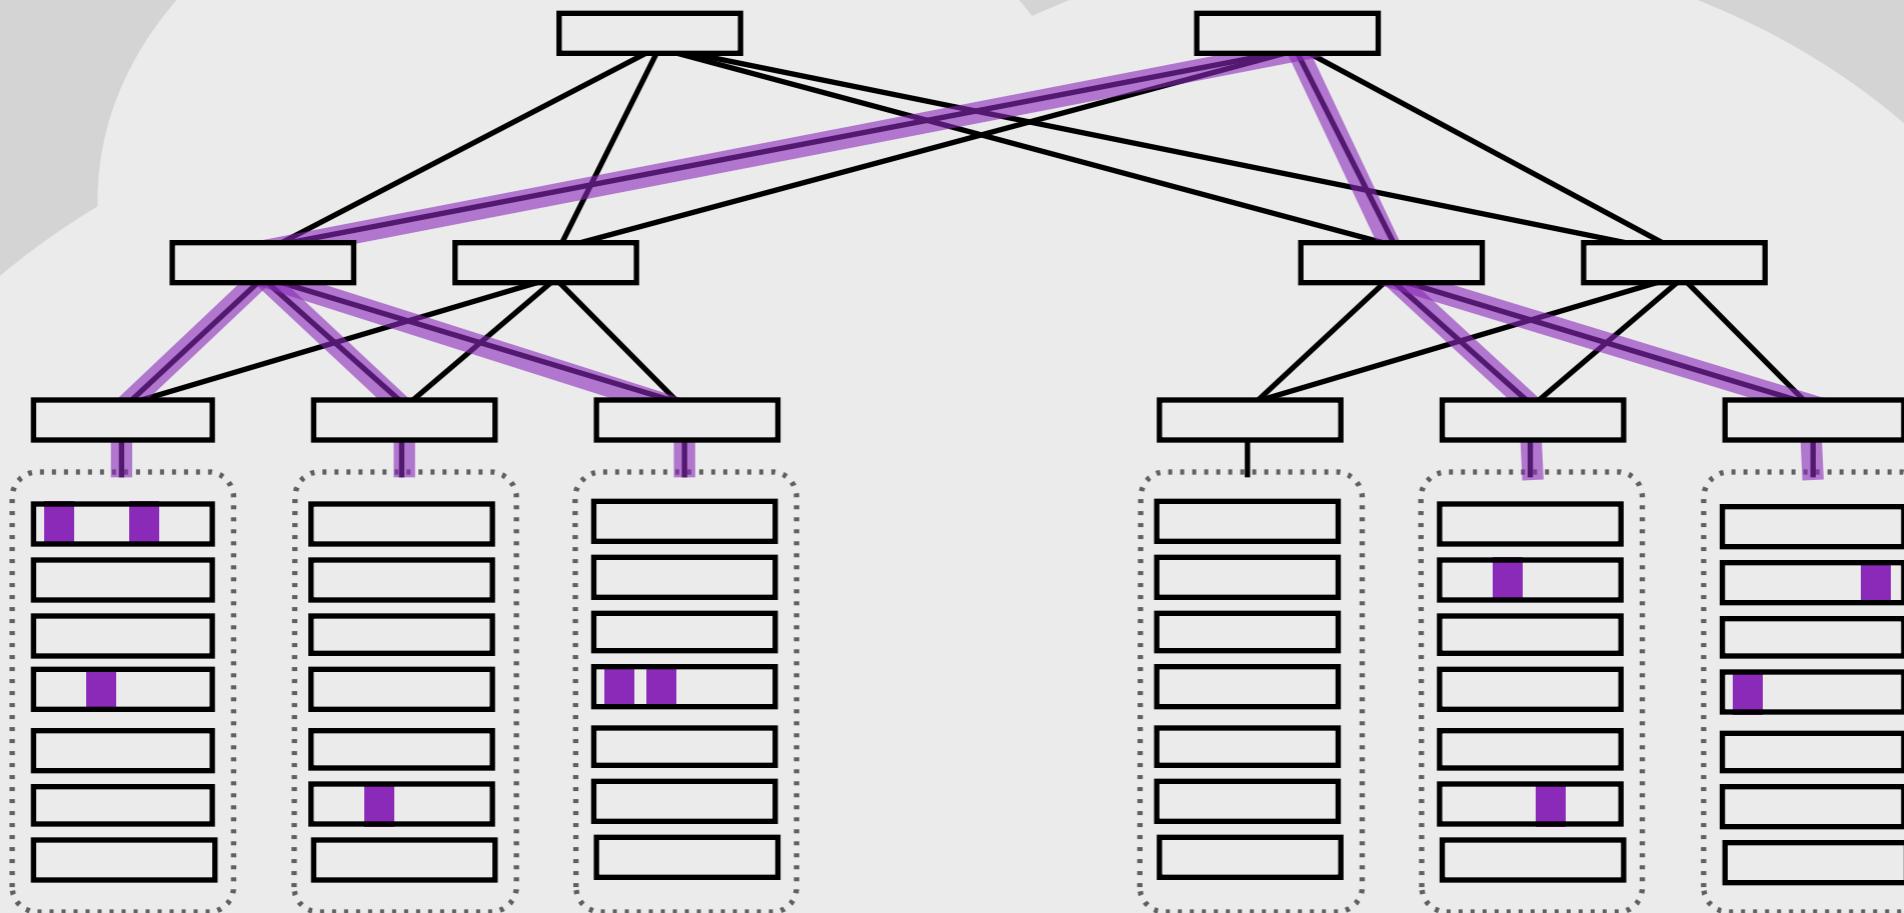
BANDWIDTH GUARANTEES



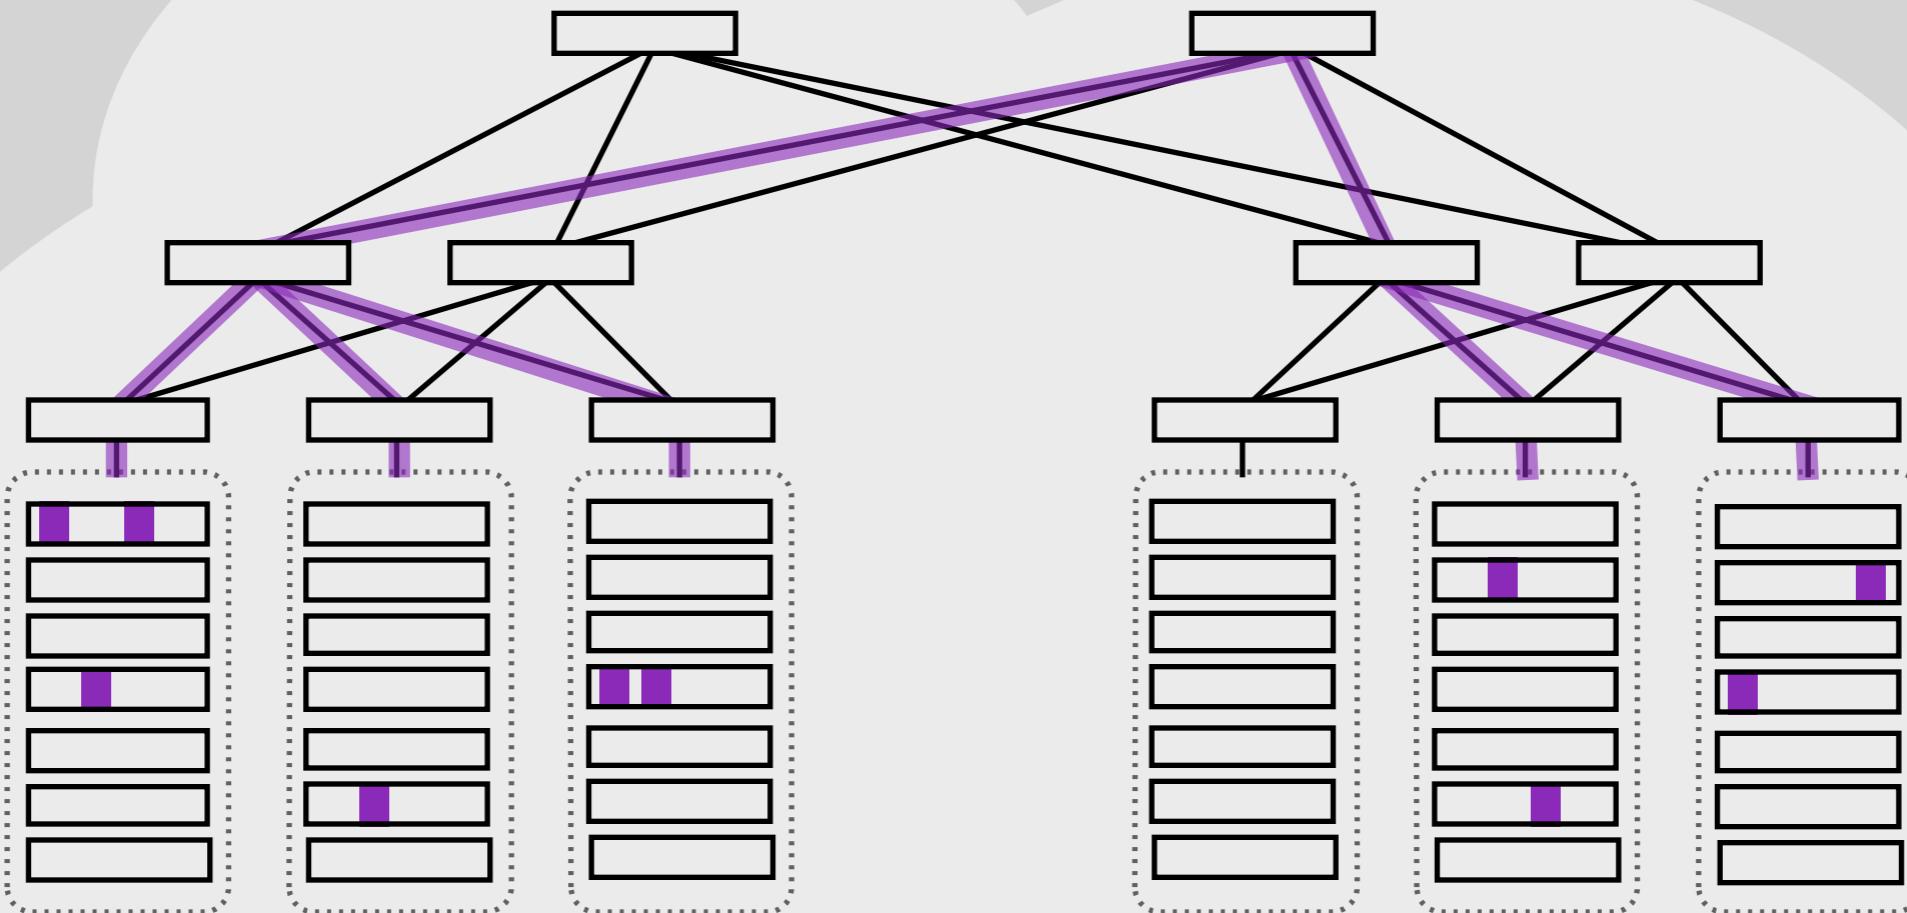
bandwidth guarantee: a promise from the **provider** to the **customer** that its VMs will be able to communicate with each other at a particular rate

(informal definition)

WHY GUARANTEES?

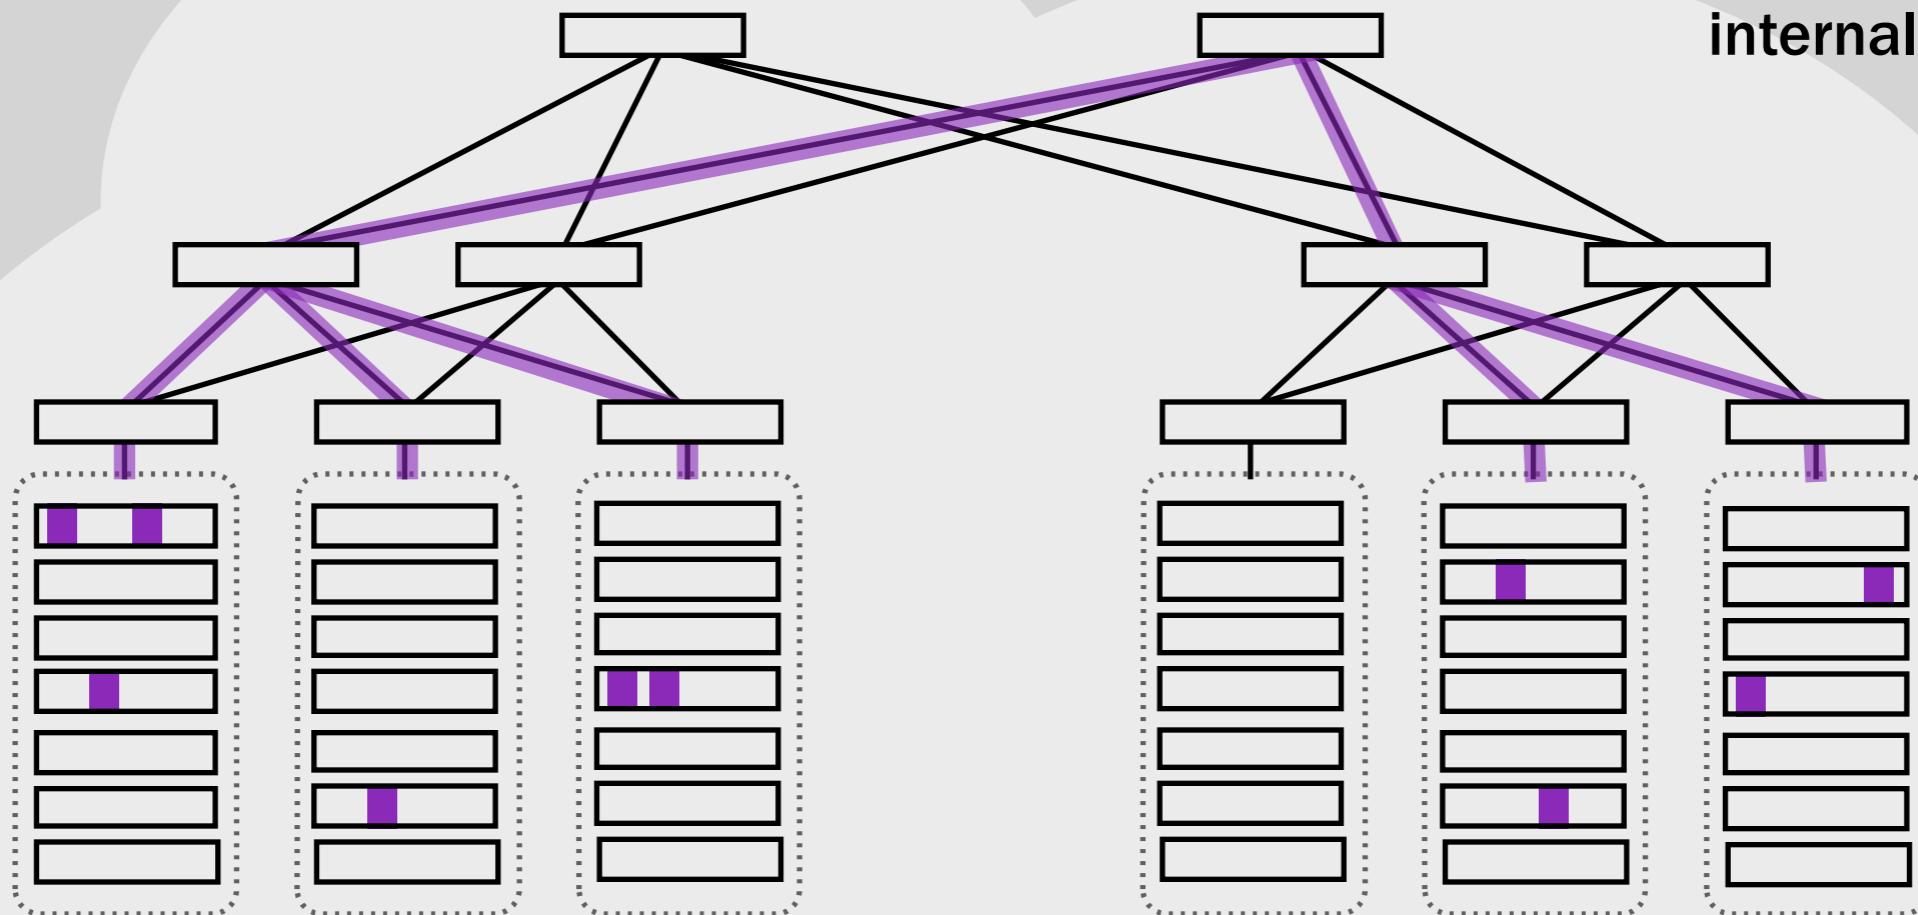


WHY GUARANTEES?



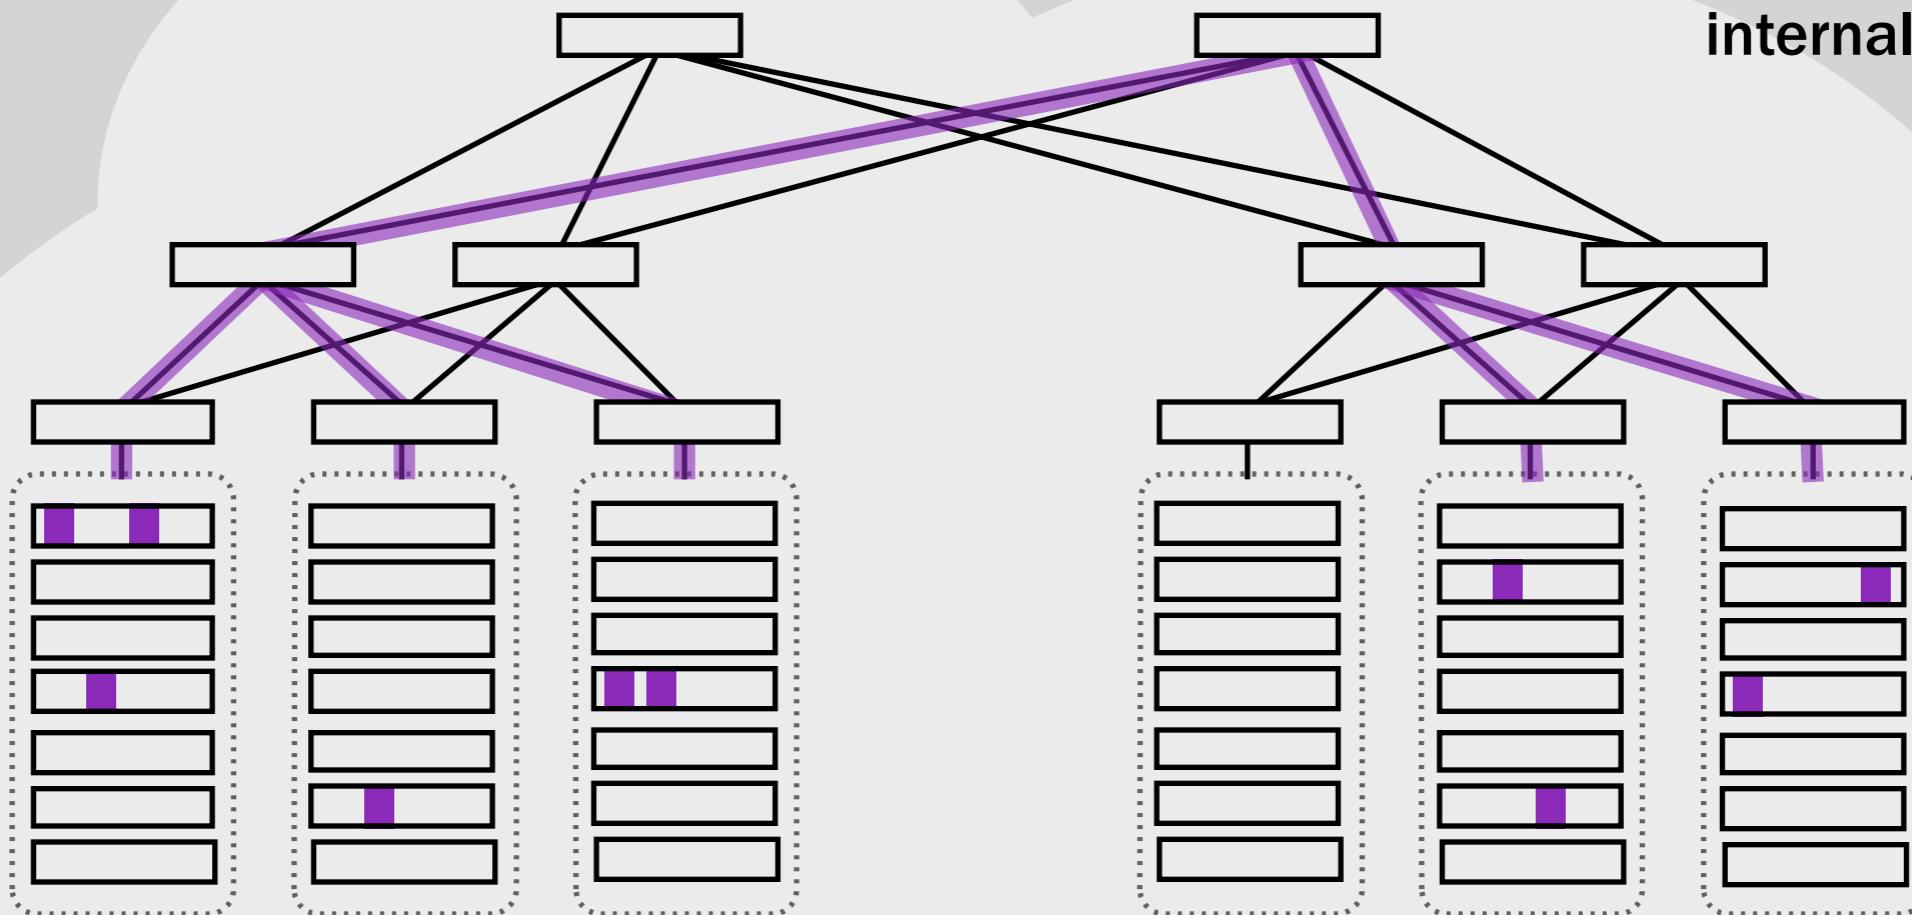
is there really that much **data**?

WHY GUARANTEES?



is there really that much **data**?

WHY GUARANTEES?

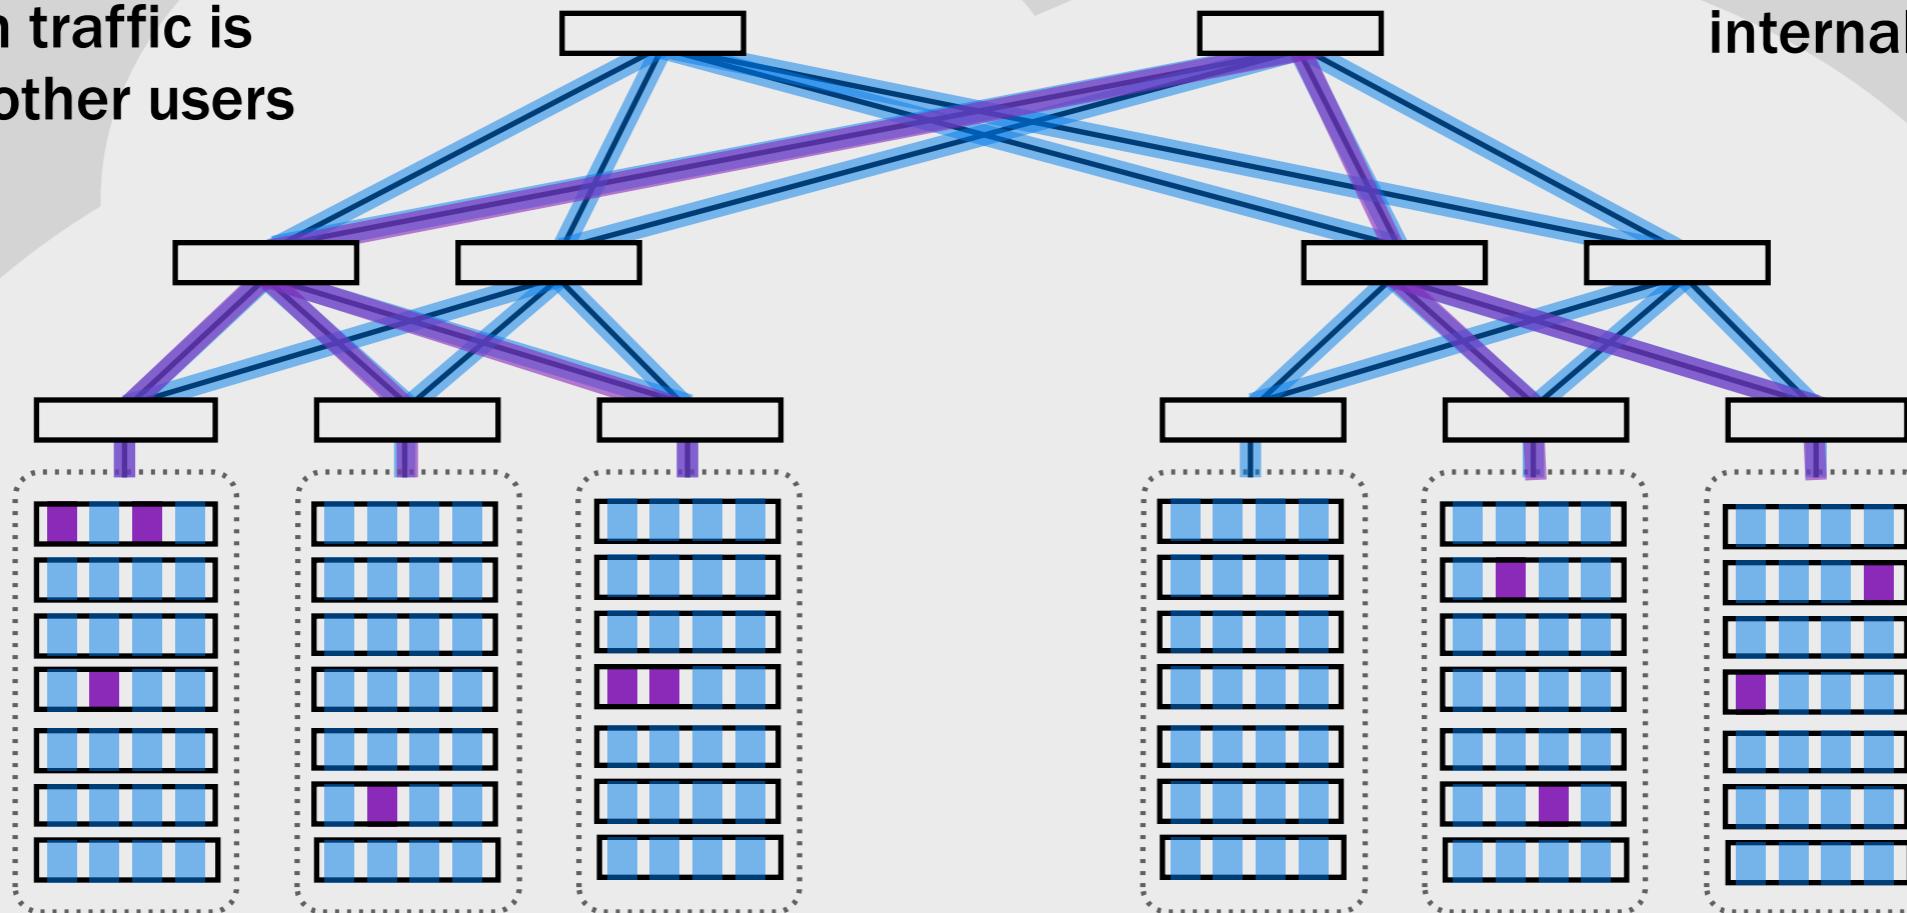


aren't cloud networks **homogeneous** and **well-provisioned**?

WHY GUARANTEES?

application traffic is affected by other users

applications can send terabytes or petabytes of data internally [1]

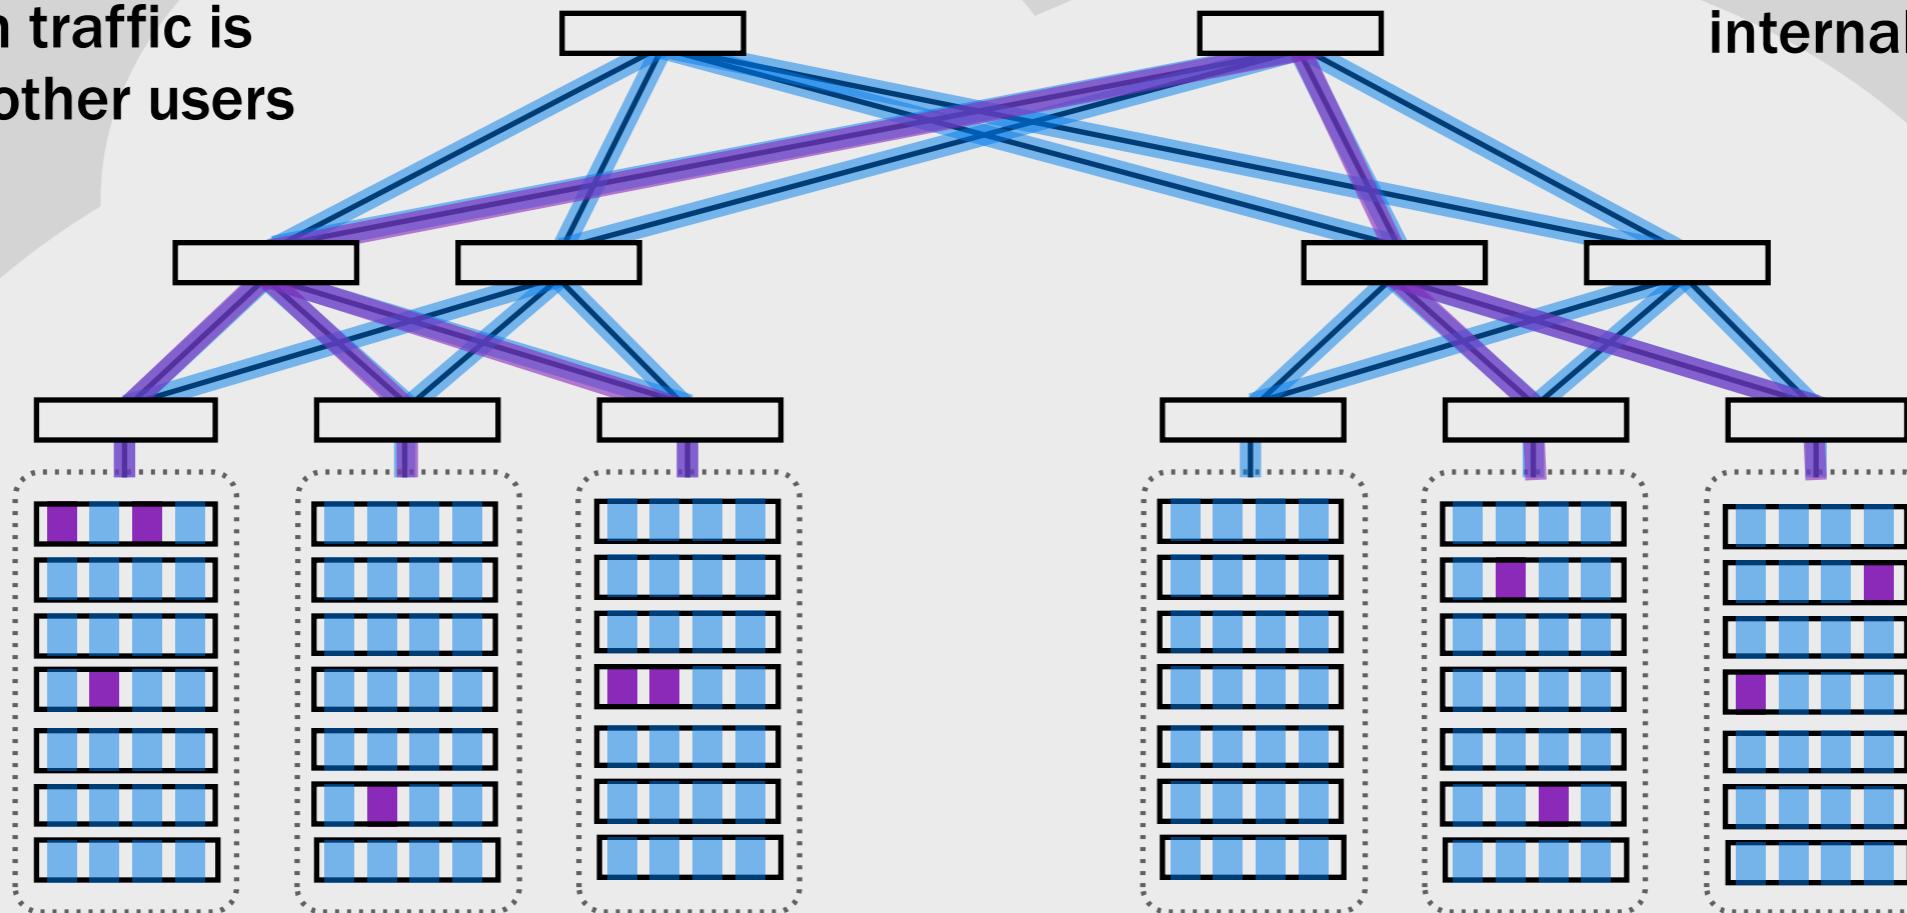


aren't cloud networks **homogeneous** and **well-provisioned**?

WHY GUARANTEES?

application traffic is affected by other users

applications can send terabytes or petabytes of data internally [1]

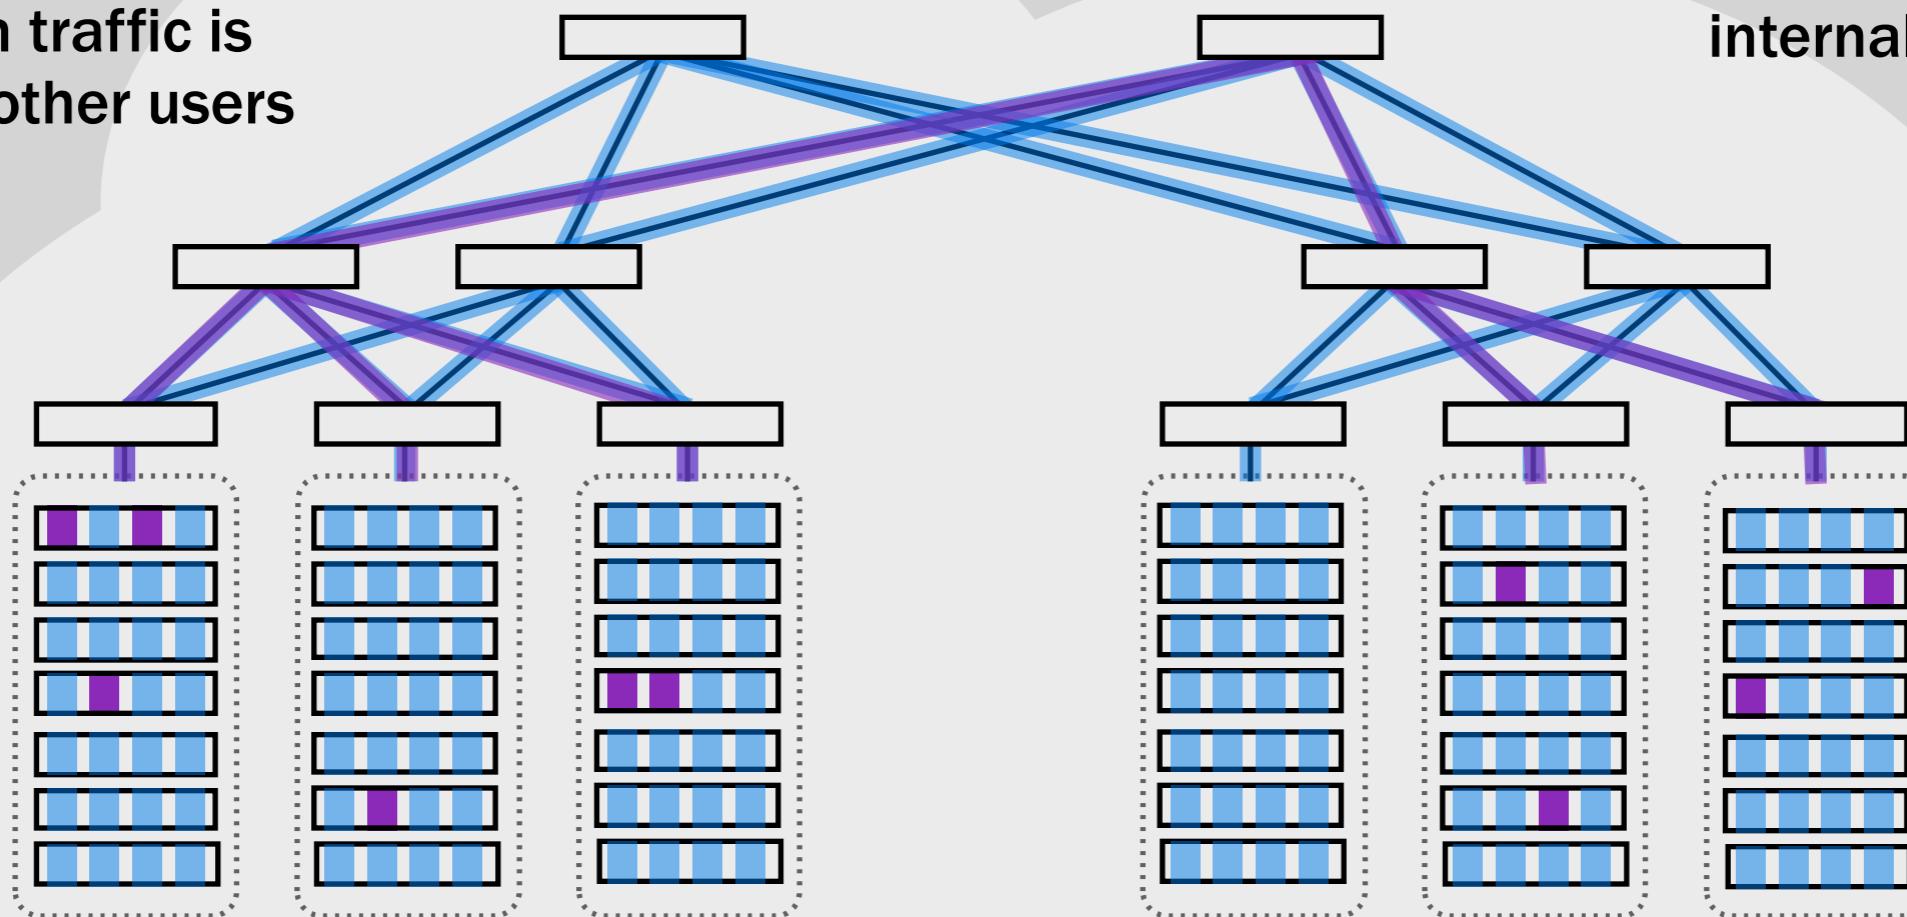


what **customers** care about network performance?

WHY GUARANTEES?

application traffic is affected by other users

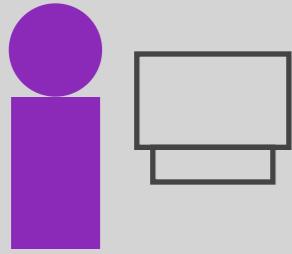
applications can send terabytes or petabytes of data internally [1]



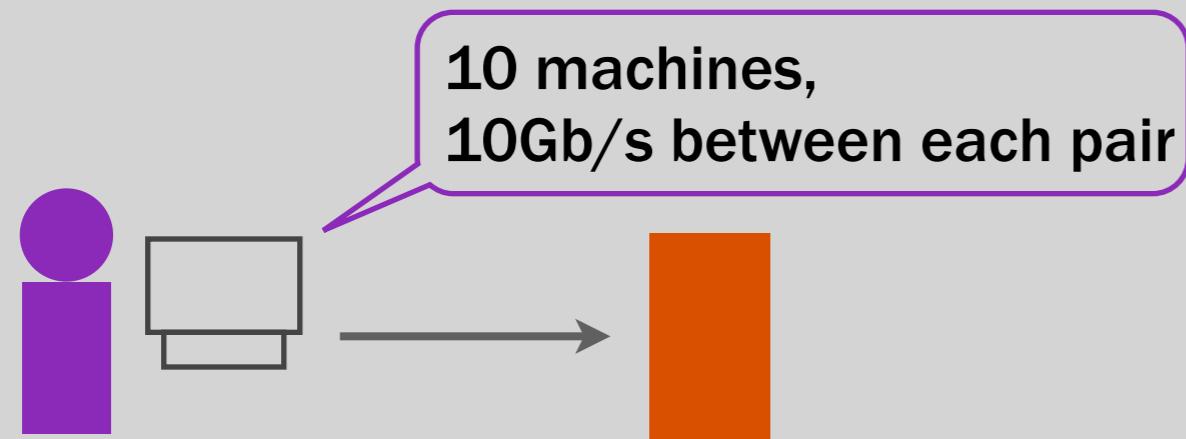
what **customers** care about network performance?

enterprise customers need to satisfy SLAs
with their own customers

POSSIBLE ARCHITECTURE

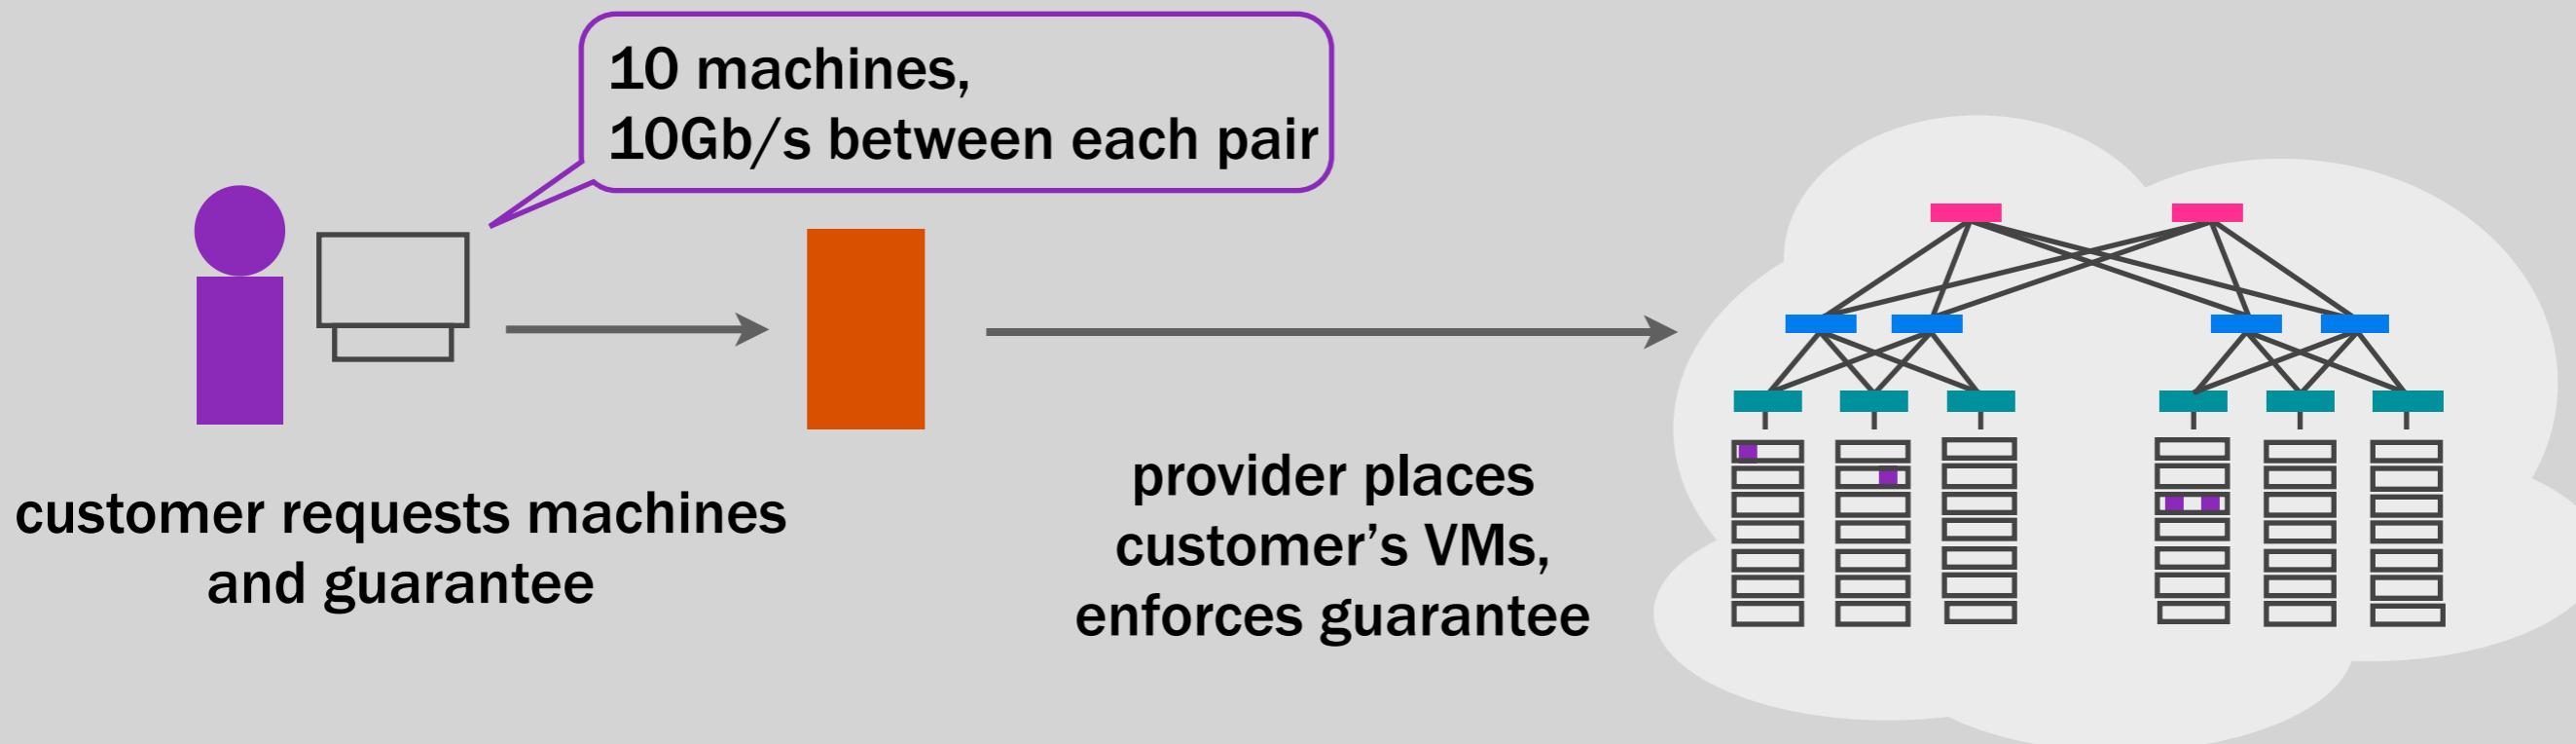


POSSIBLE ARCHITECTURE

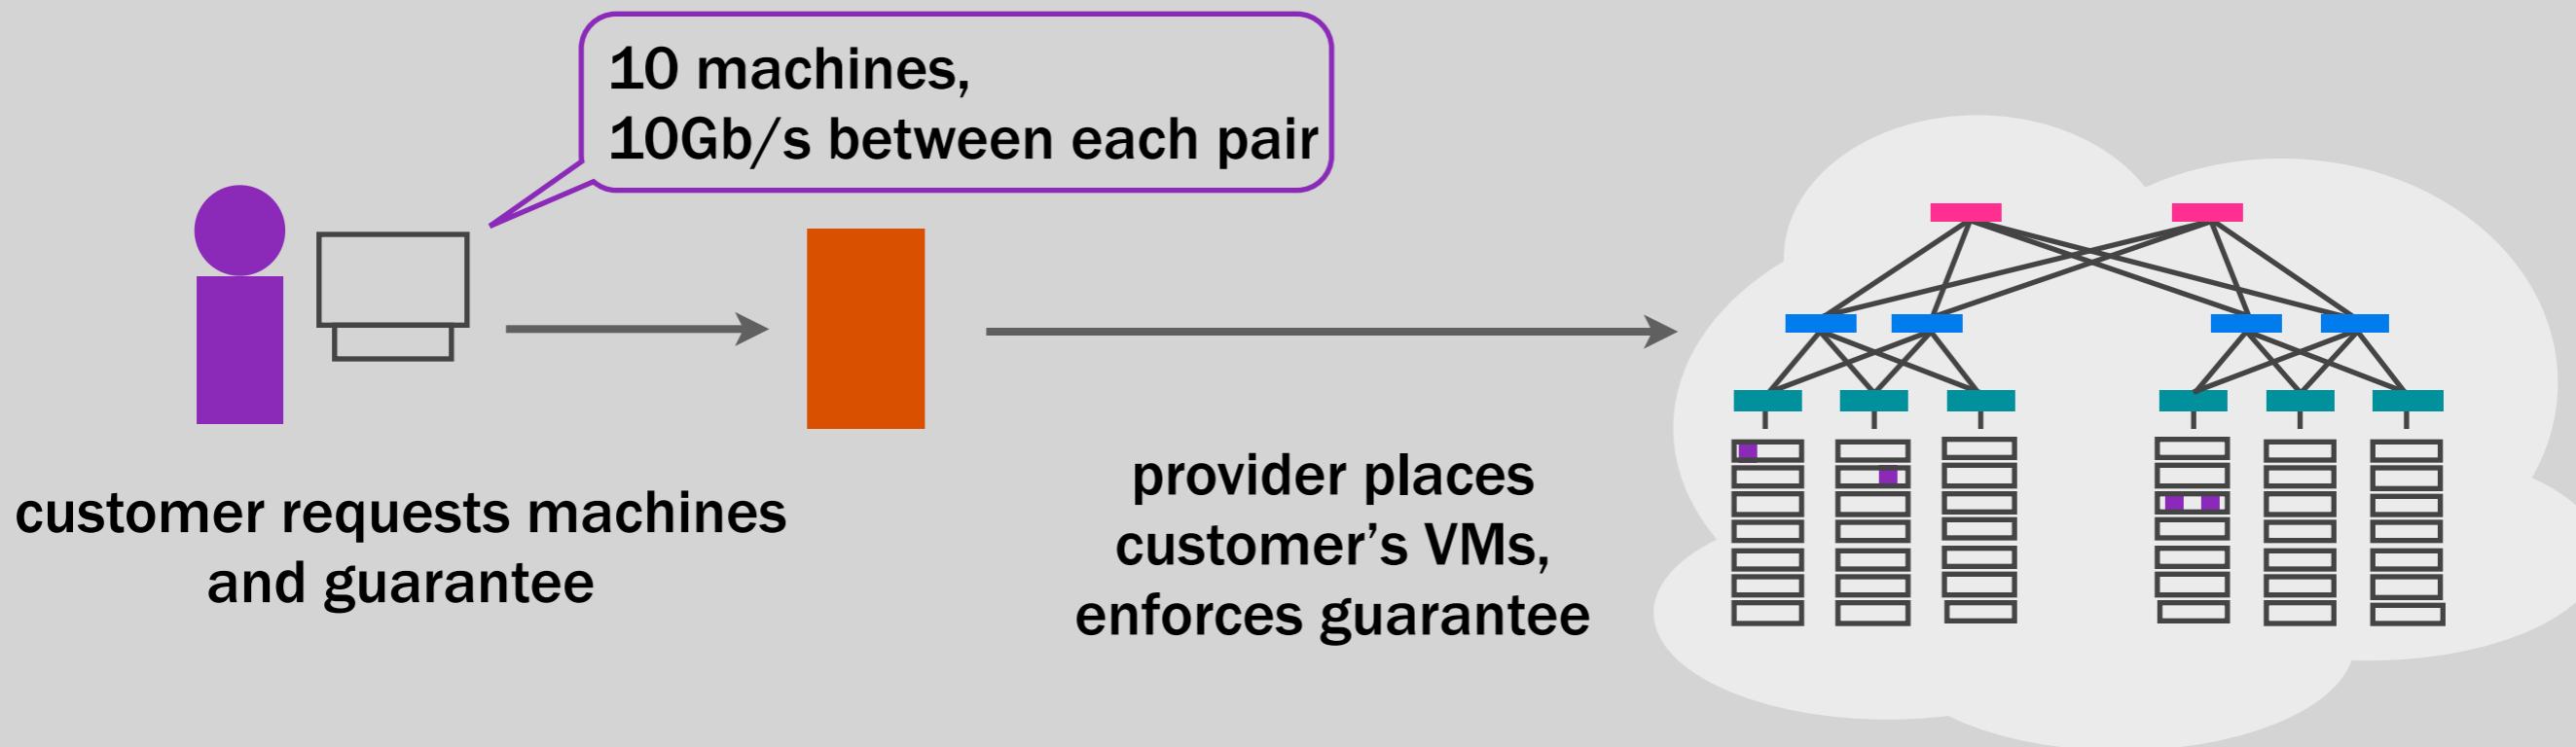


customer requests machines
and guarantee

POSSIBLE ARCHITECTURE

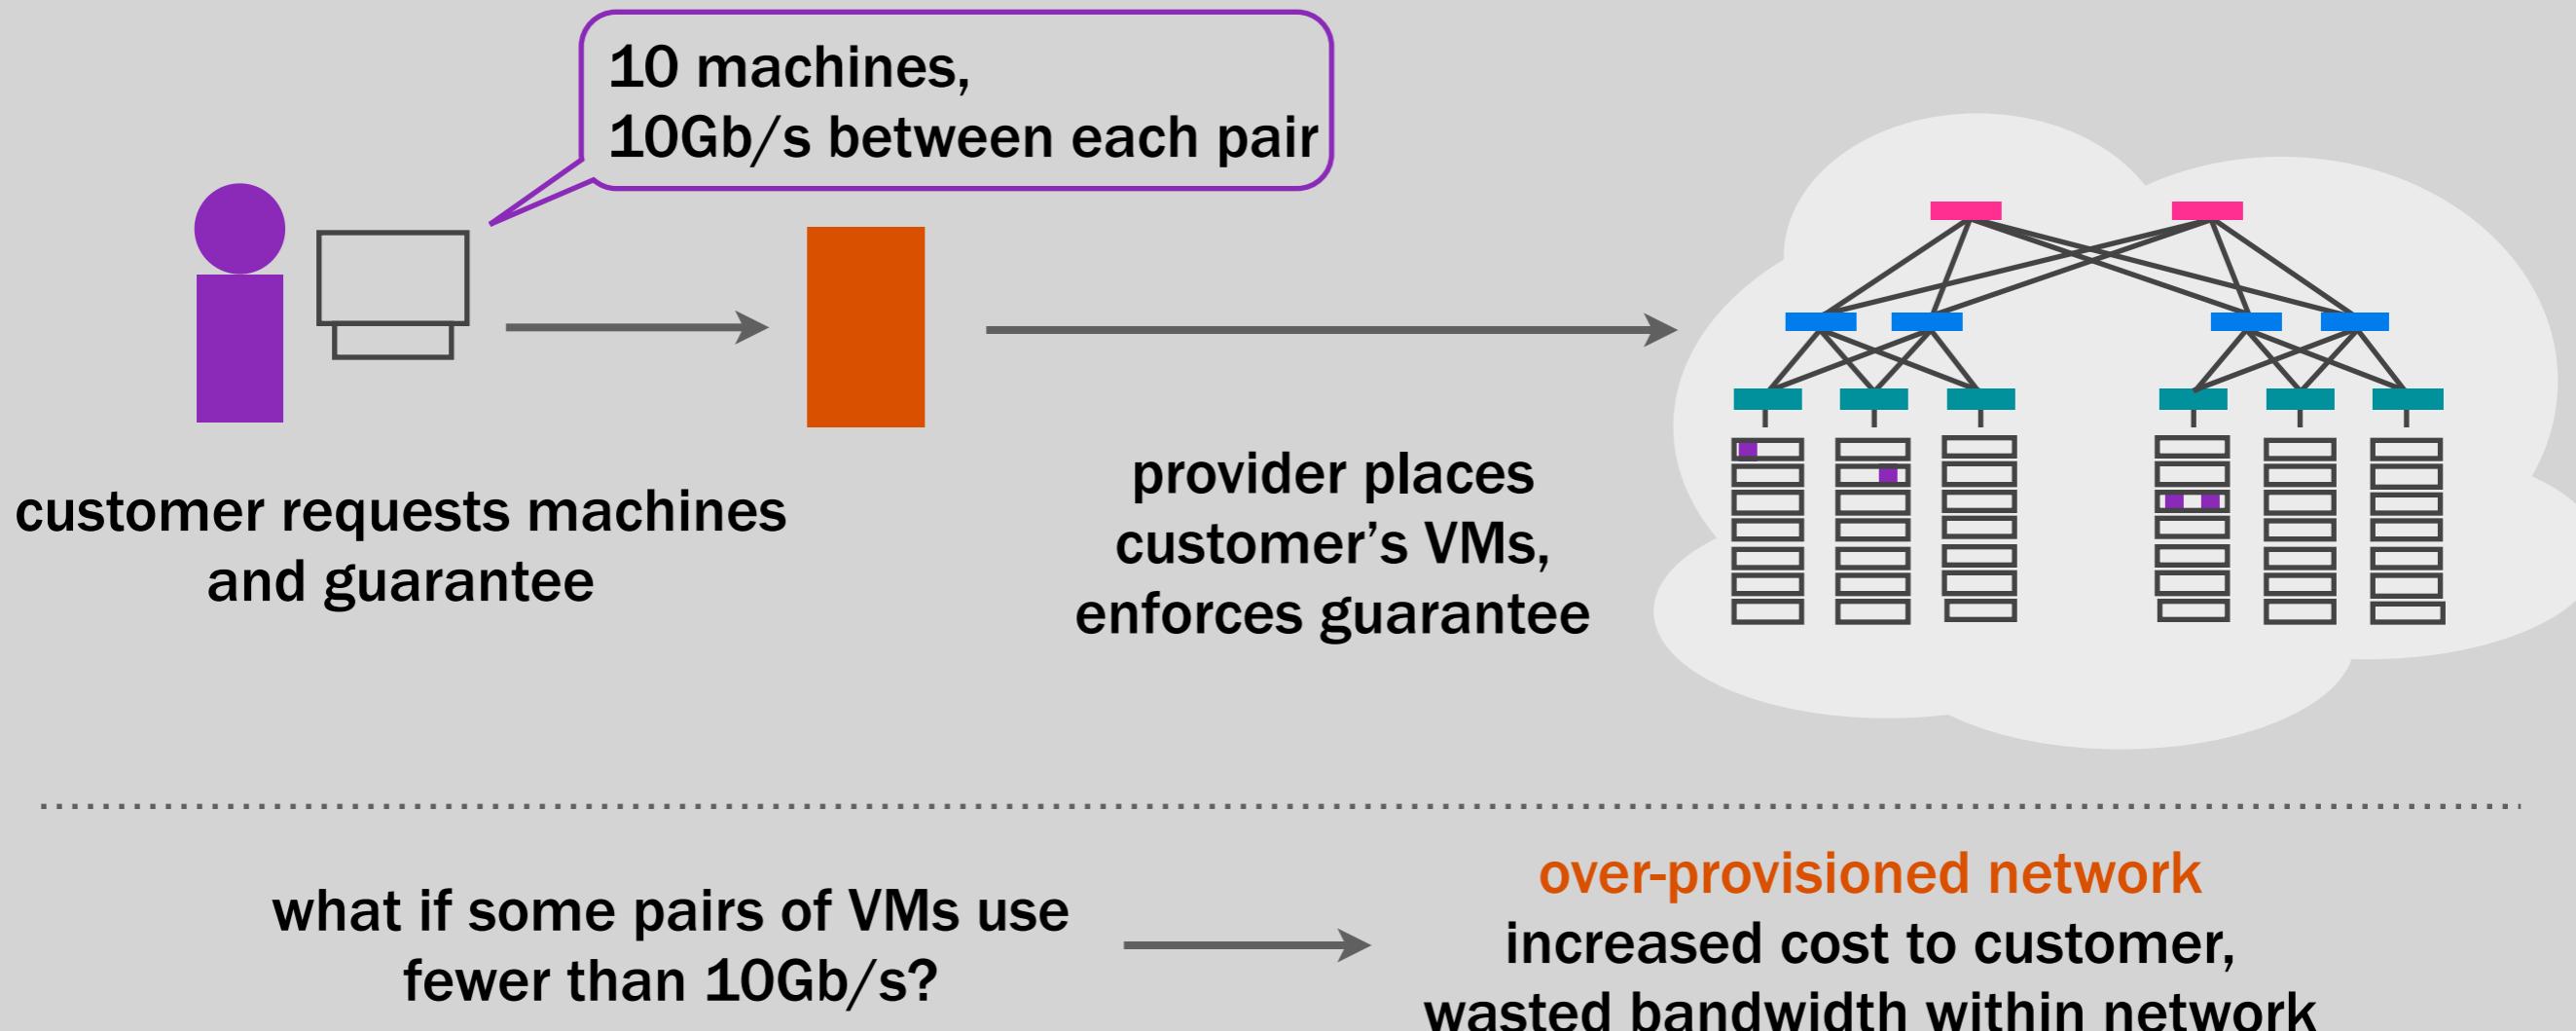


POSSIBLE ARCHITECTURE

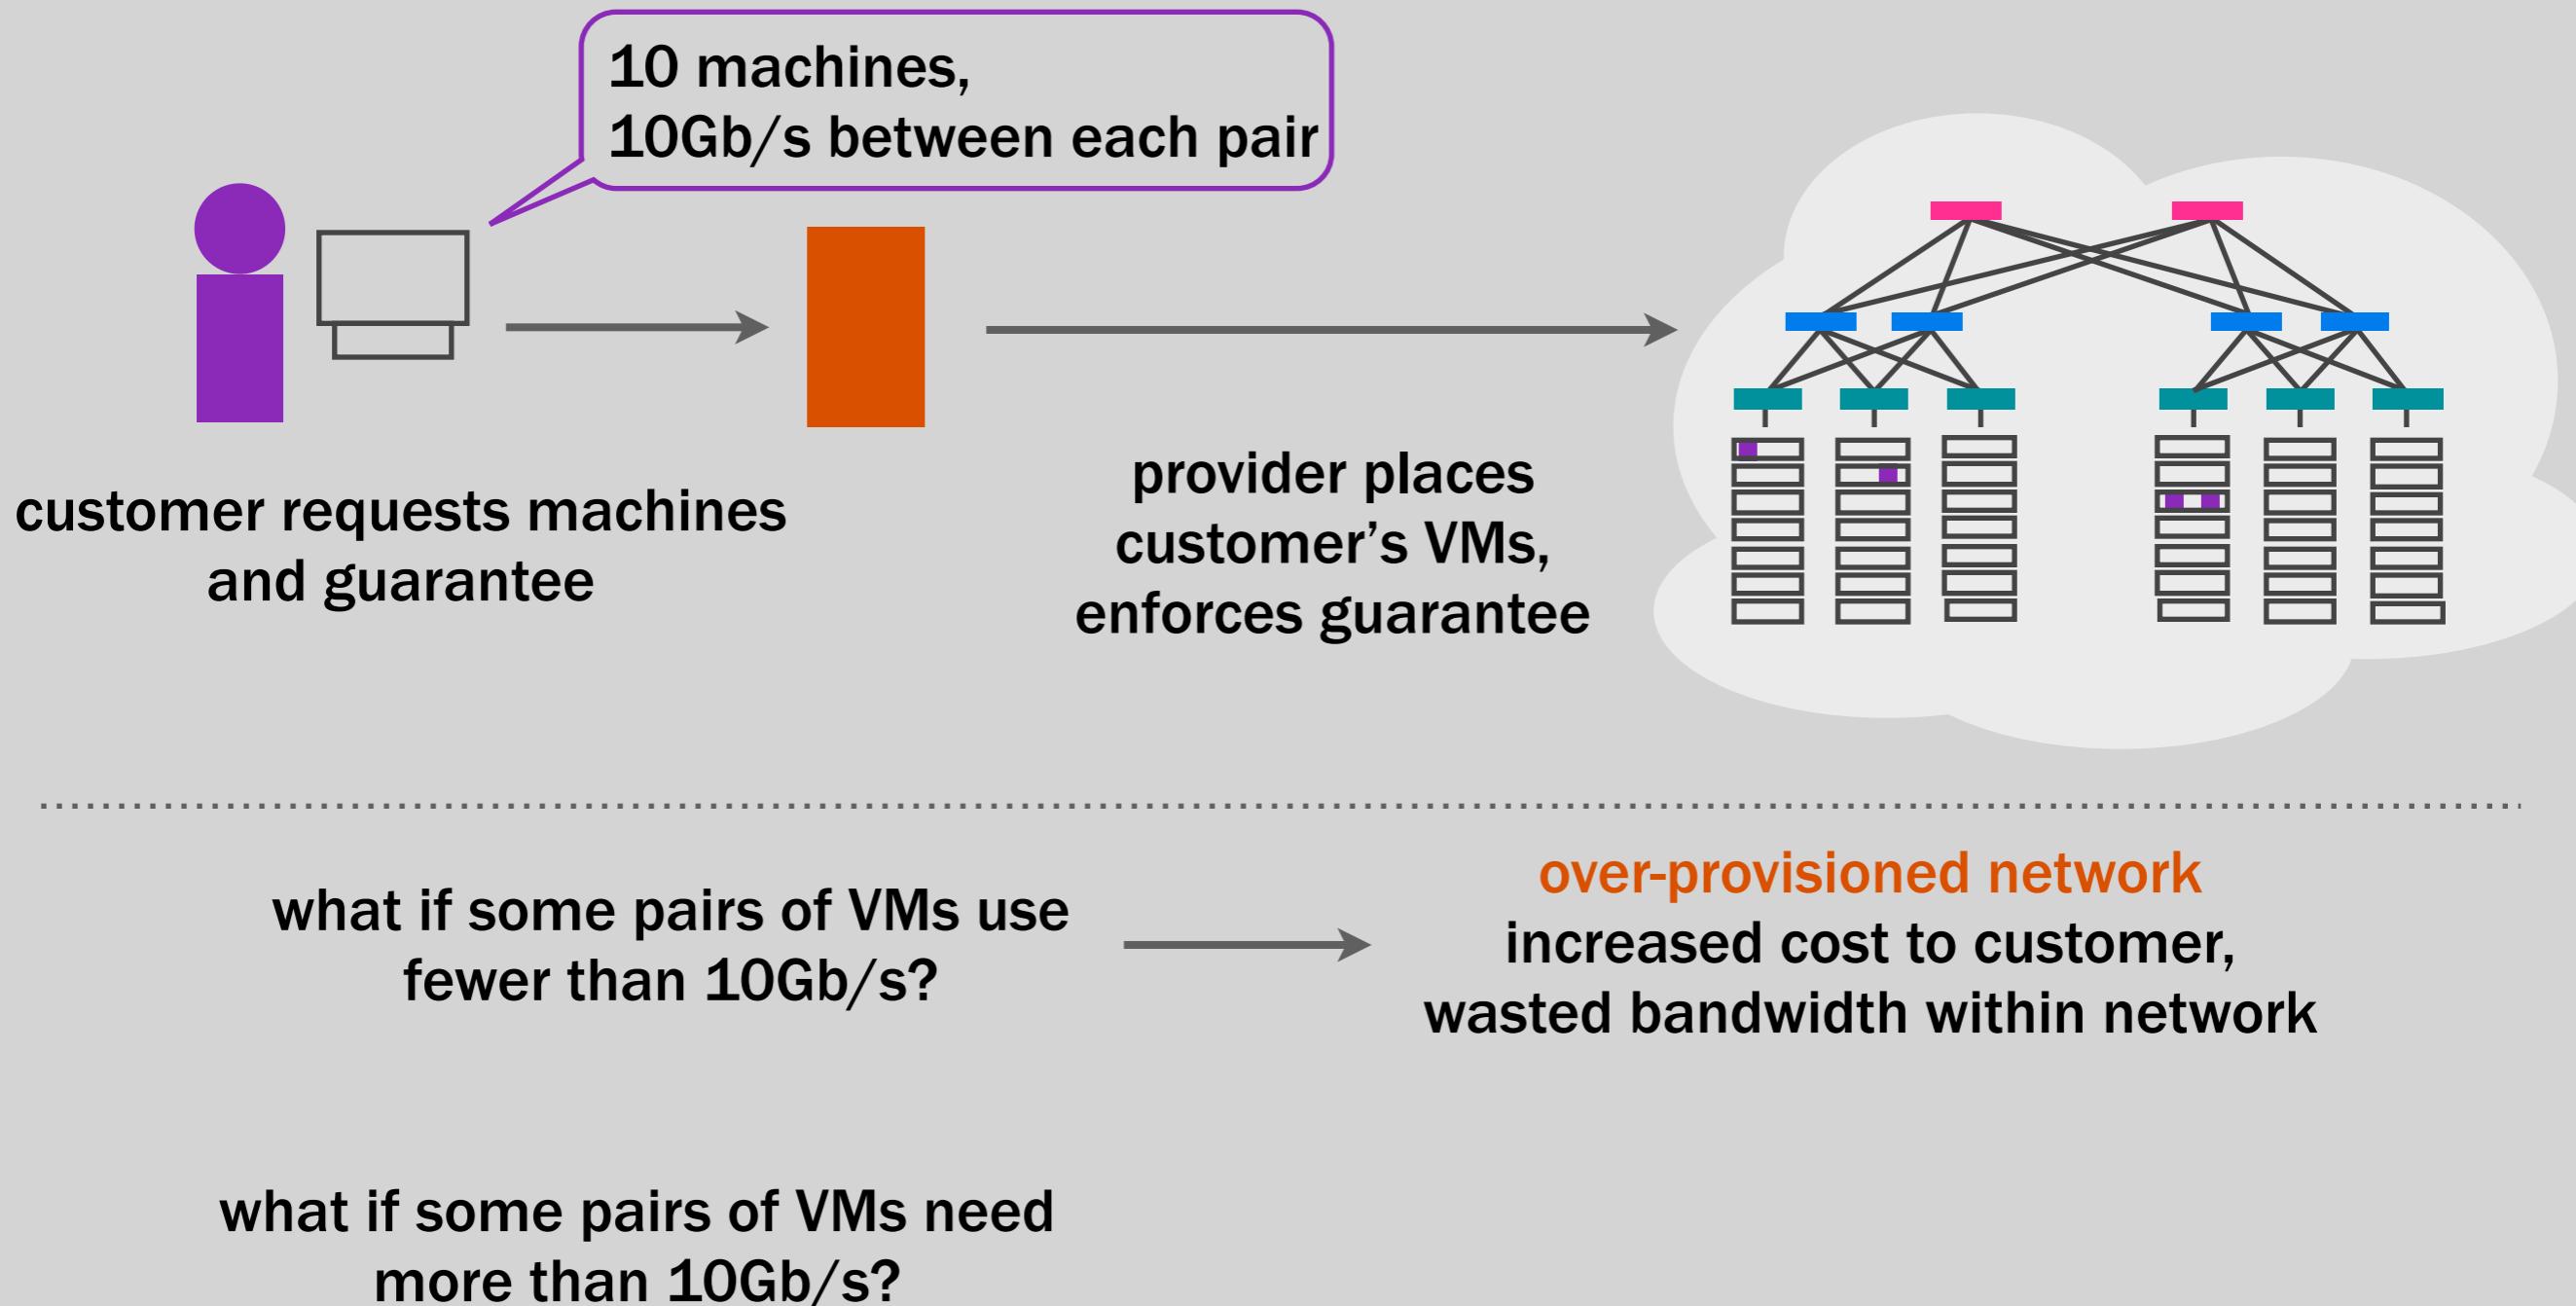


what if some pairs of VMs use fewer than 10Gb/s?

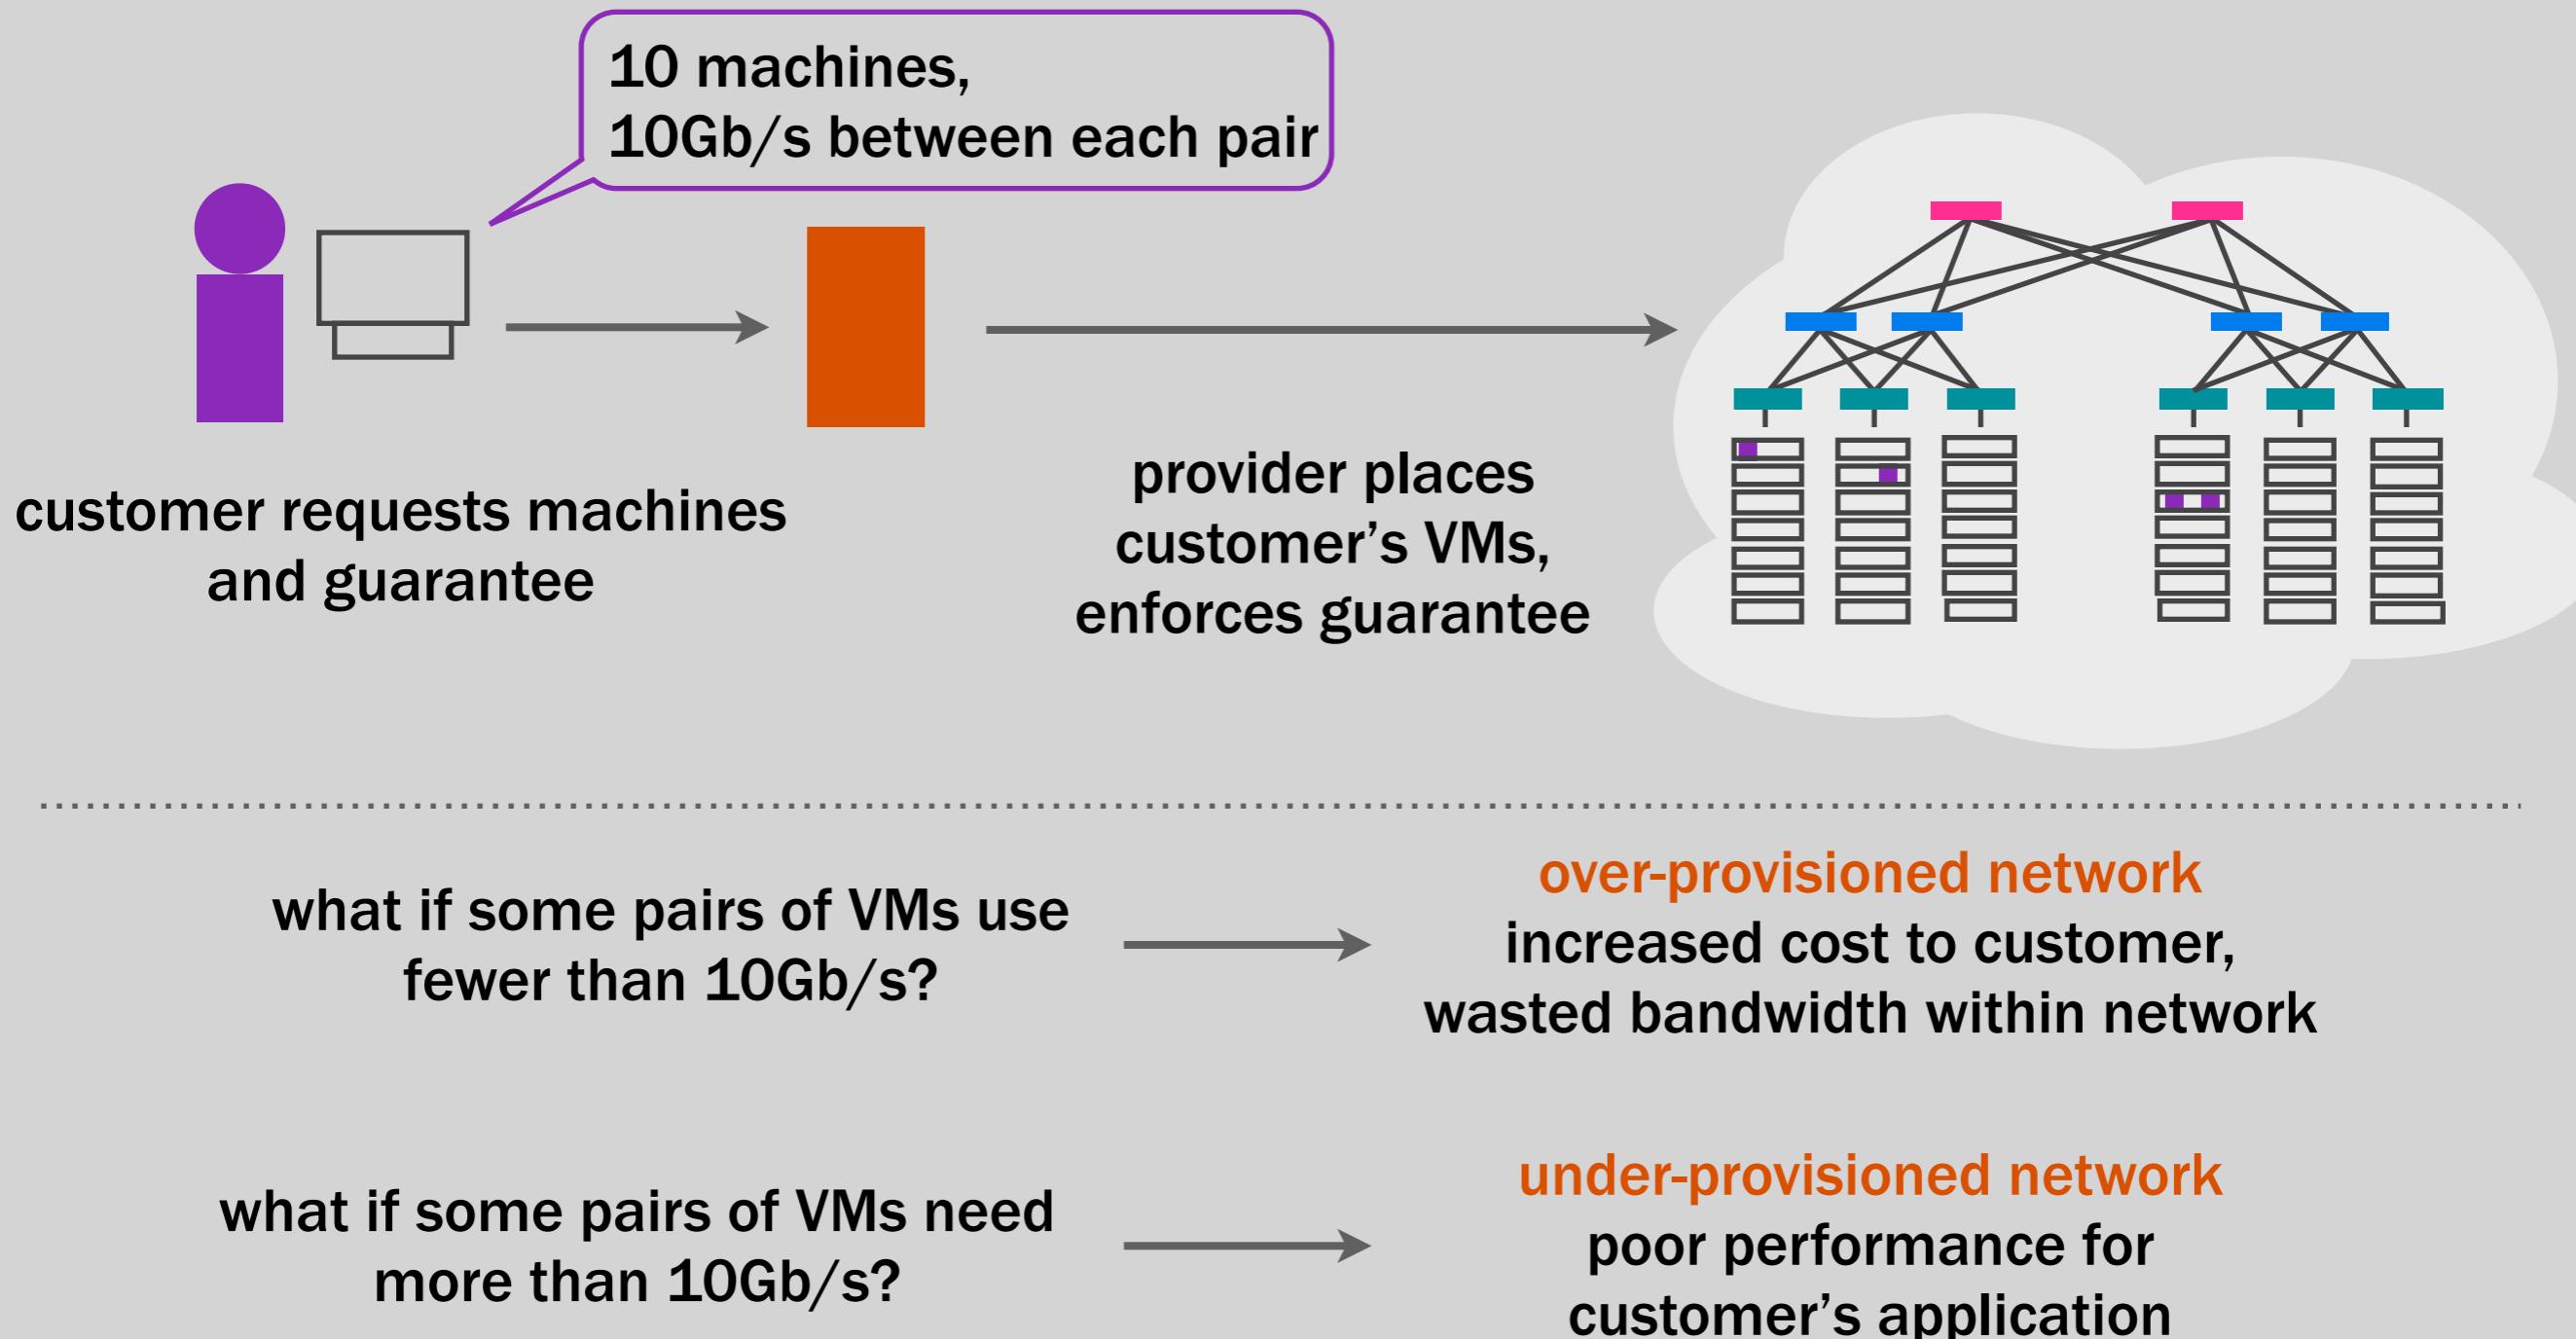
POSSIBLE ARCHITECTURE



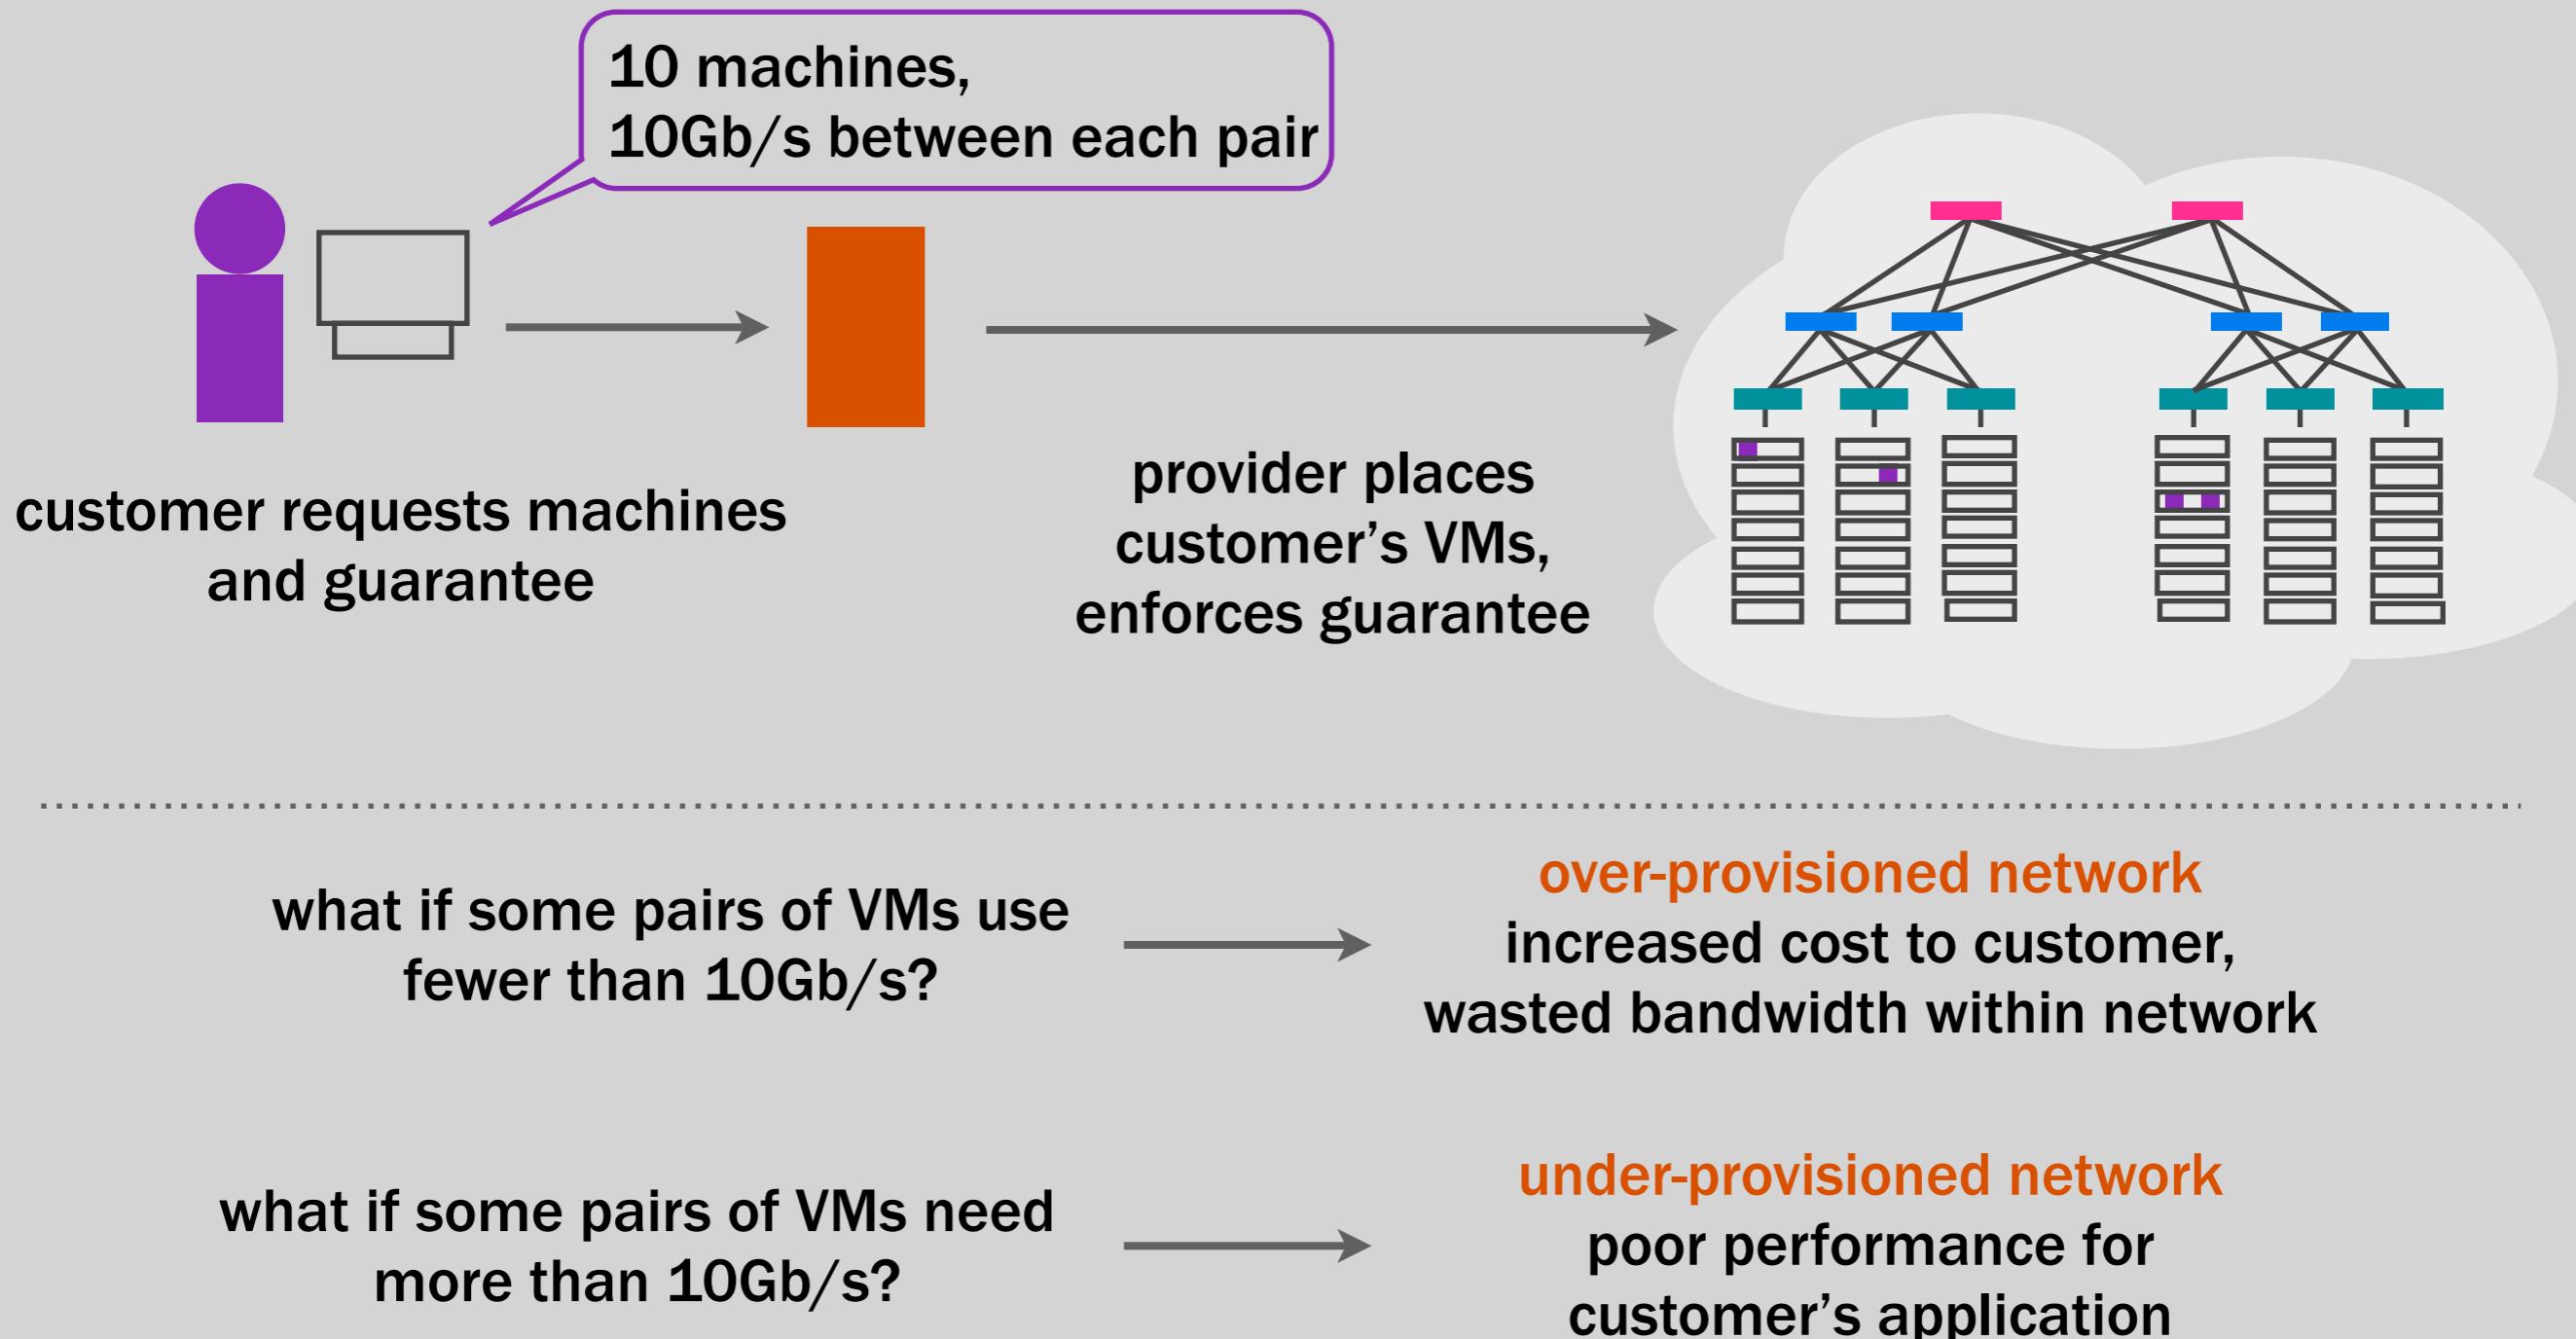
POSSIBLE ARCHITECTURE



POSSIBLE ARCHITECTURE



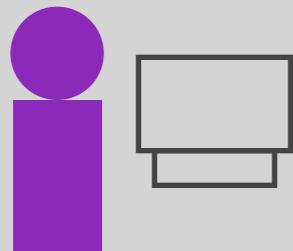
POSSIBLE ARCHITECTURE



problem: how do customers know what guarantee they need?

CICADA'S ARCHITECTURE

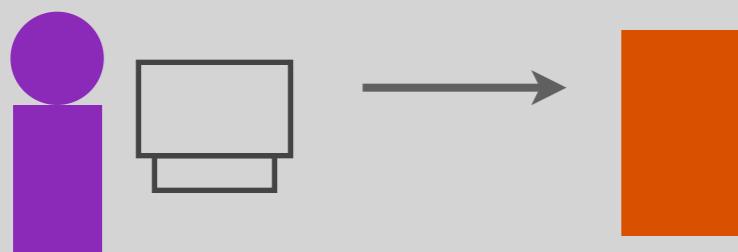
cicada makes **predictions** about an application's traffic to automatically generate a guarantee



CICADA'S ARCHITECTURE

cicada makes **predictions** about an application's traffic to automatically generate a guarantee

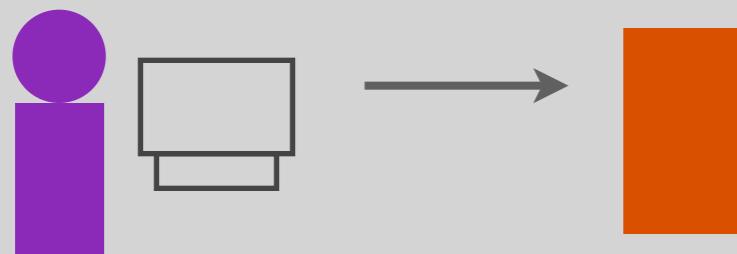
customer makes initial request for machines



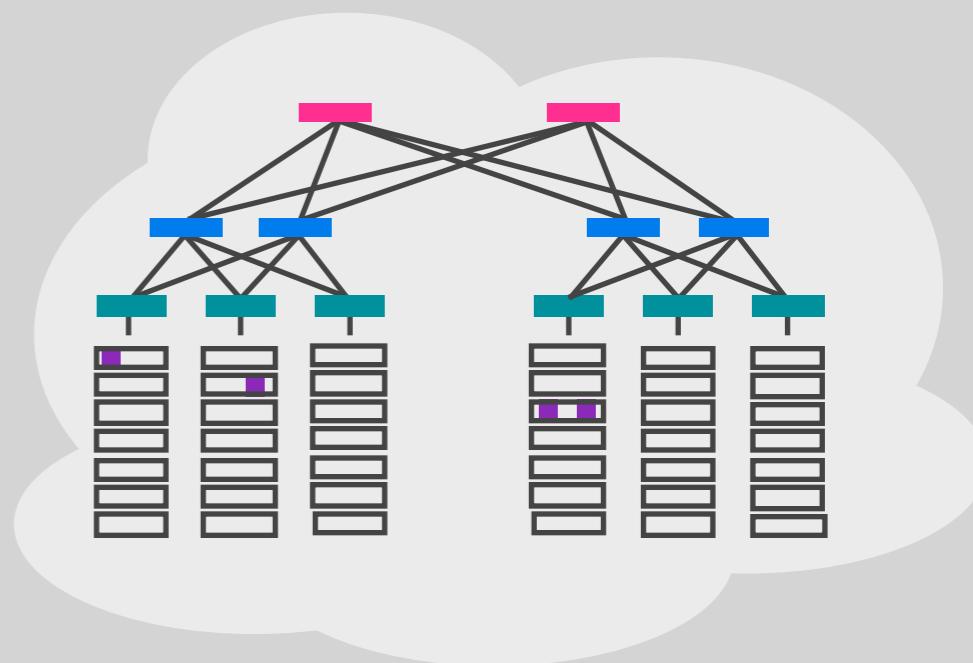
CICADA'S ARCHITECTURE

cicada makes **predictions** about an application's traffic to automatically generate a guarantee

customer makes initial request for machines



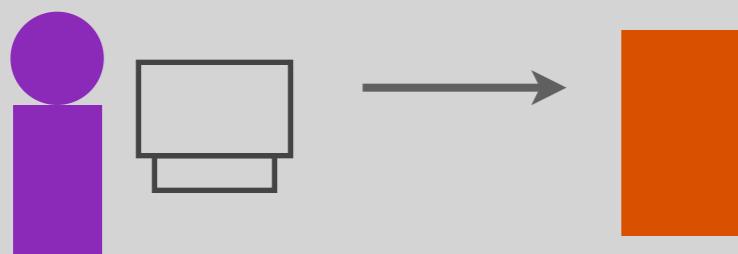
provider places tenant VMs



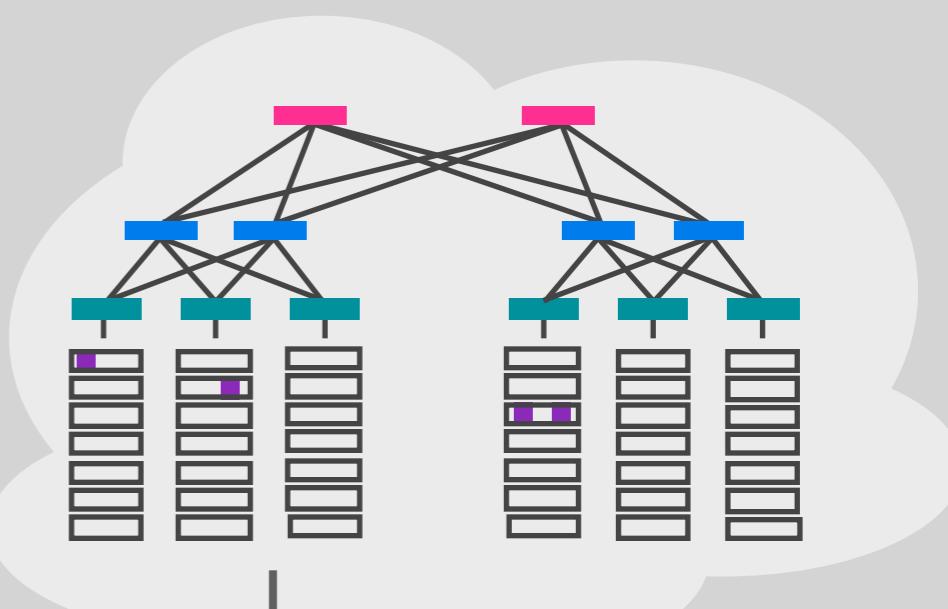
CICADA'S ARCHITECTURE

cicada makes **predictions** about an application's traffic to automatically generate a guarantee

customer makes initial request for machines



provider places tenant VMs



hypervisors send measurements to cicada controller

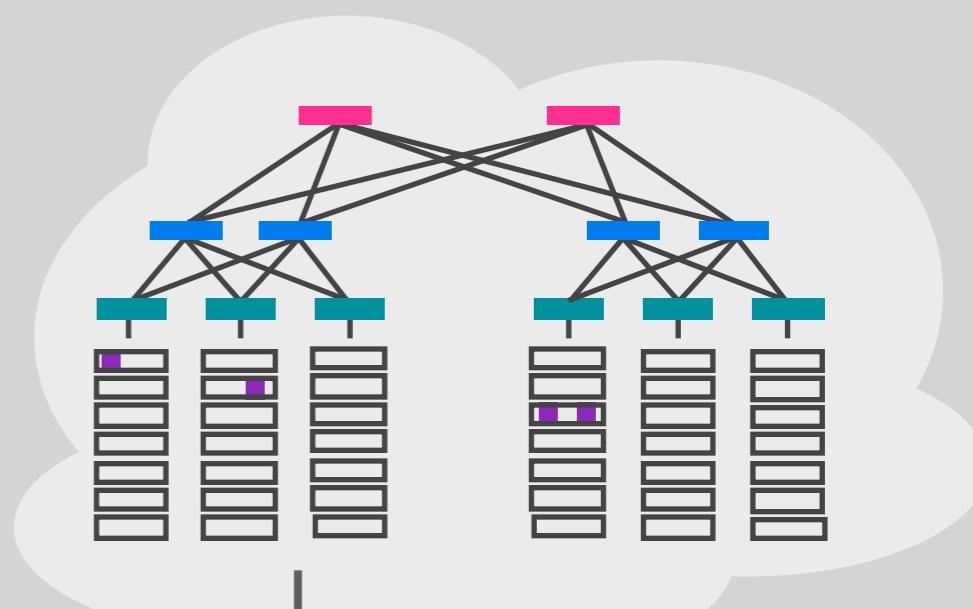
CICADA'S ARCHITECTURE

cicada makes **predictions** about an application's traffic to automatically generate a guarantee

customer makes initial request for machines



provider places tenant VMs



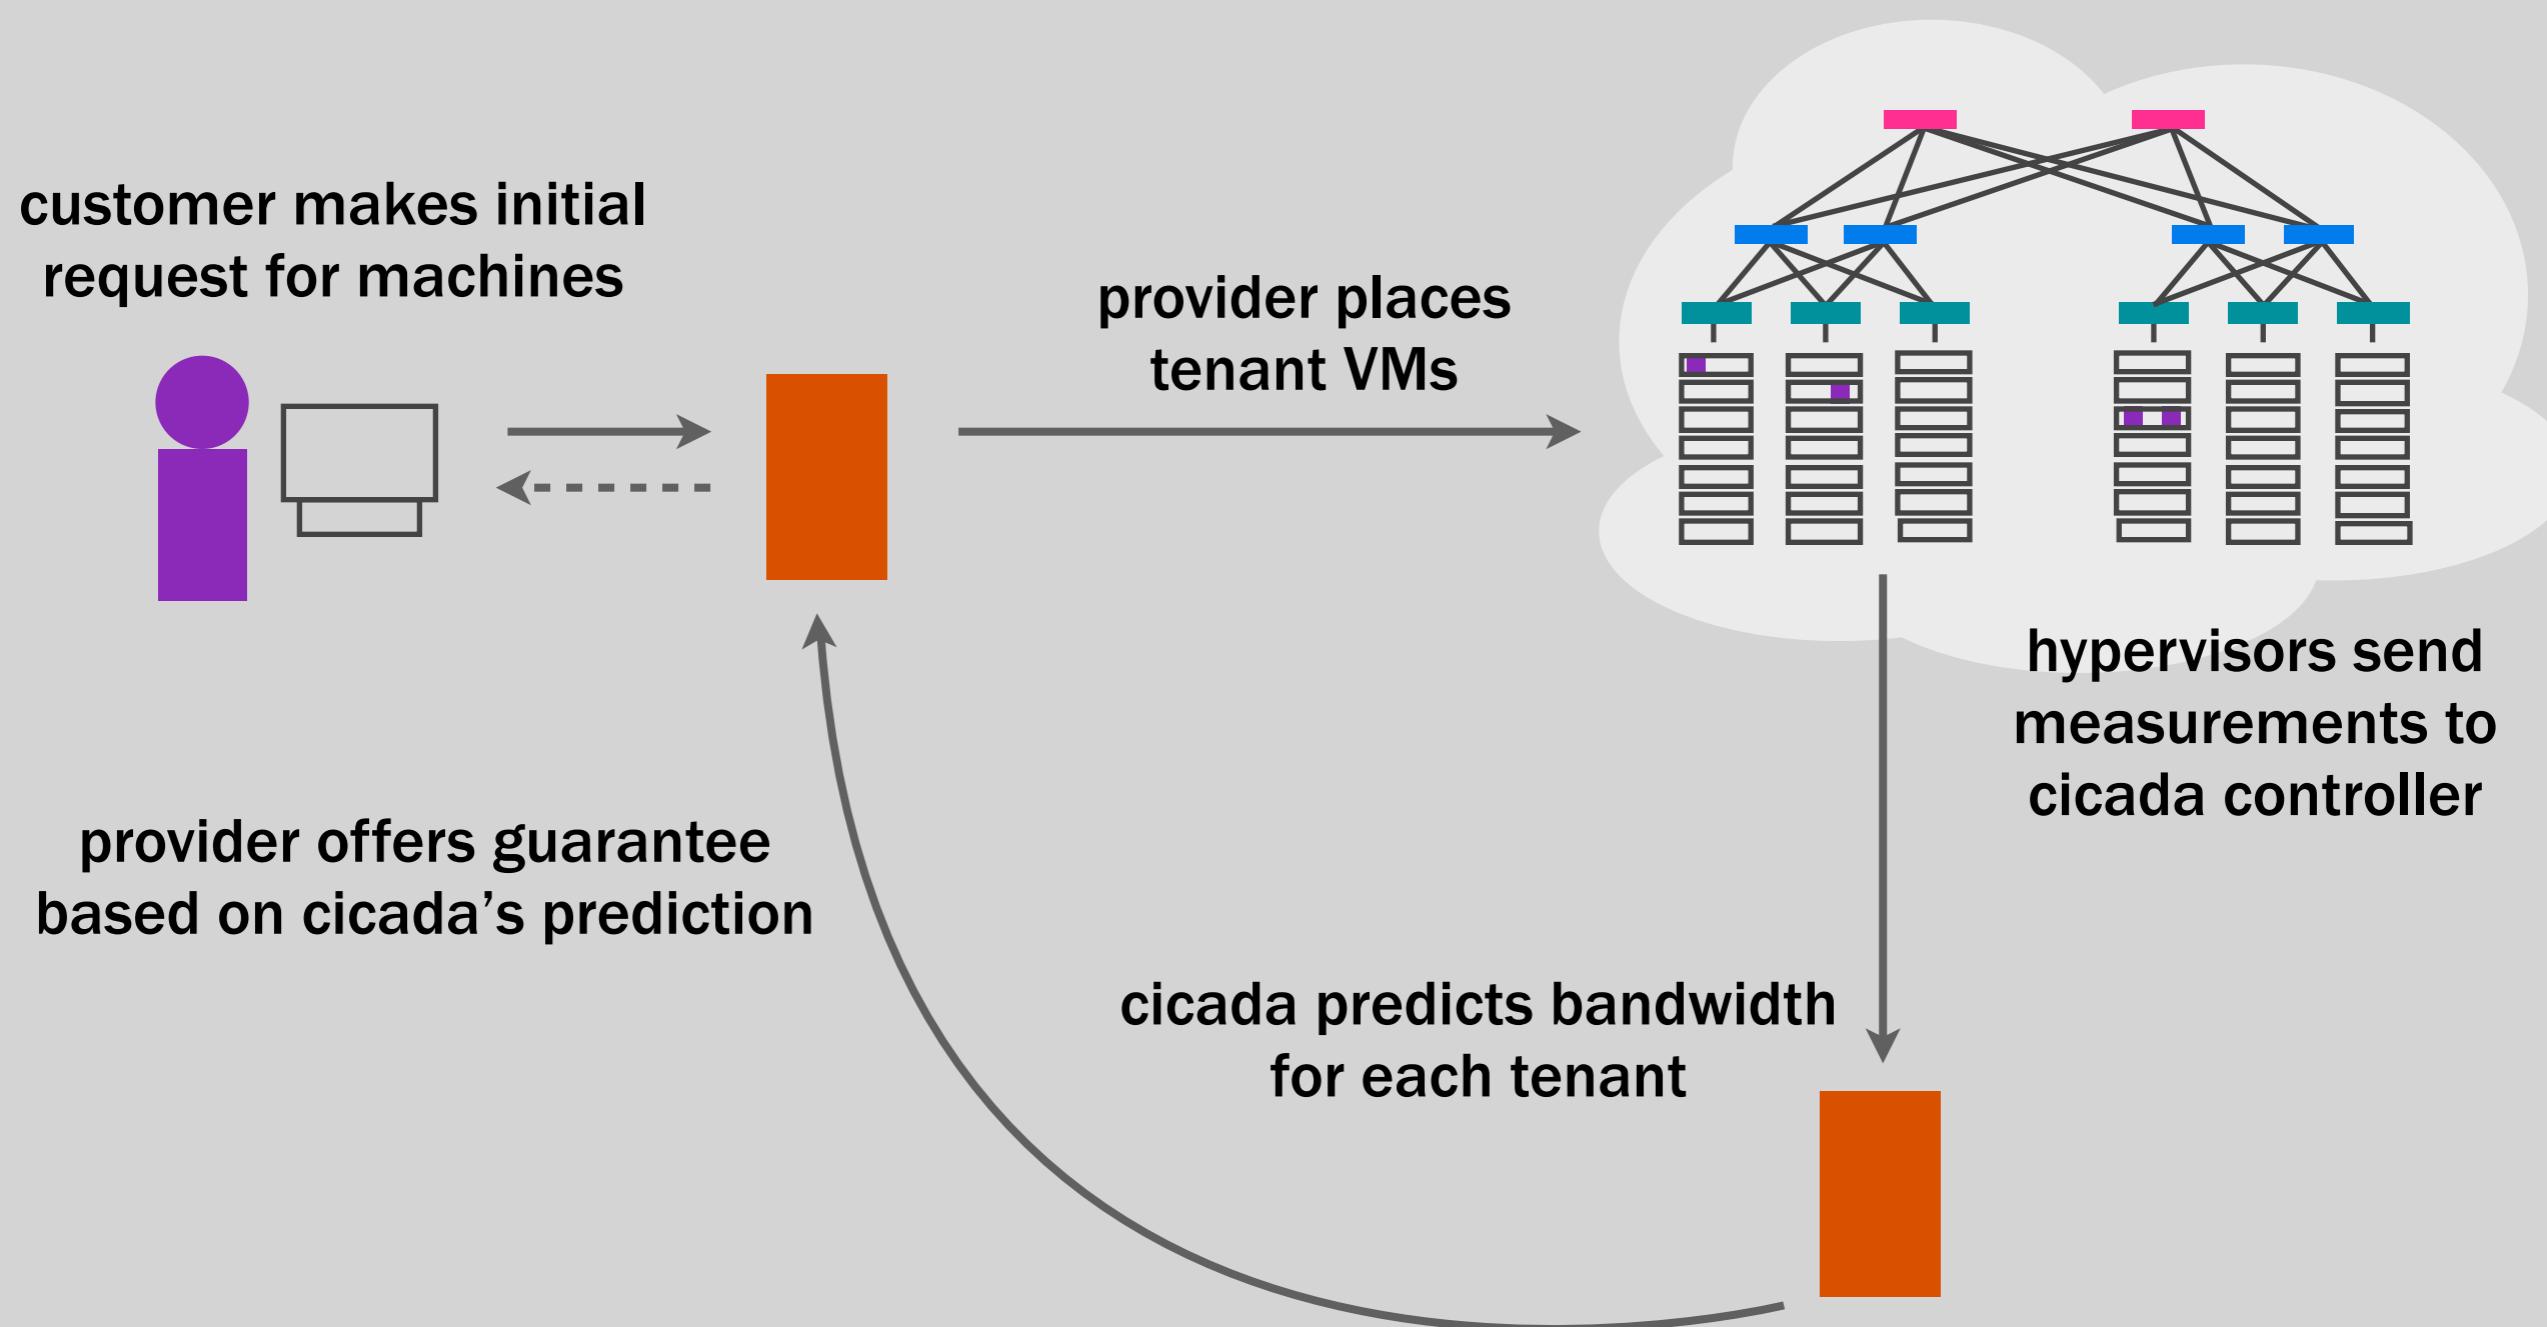
hypervisors send measurements to cicada controller

cicada predicts bandwidth for each tenant



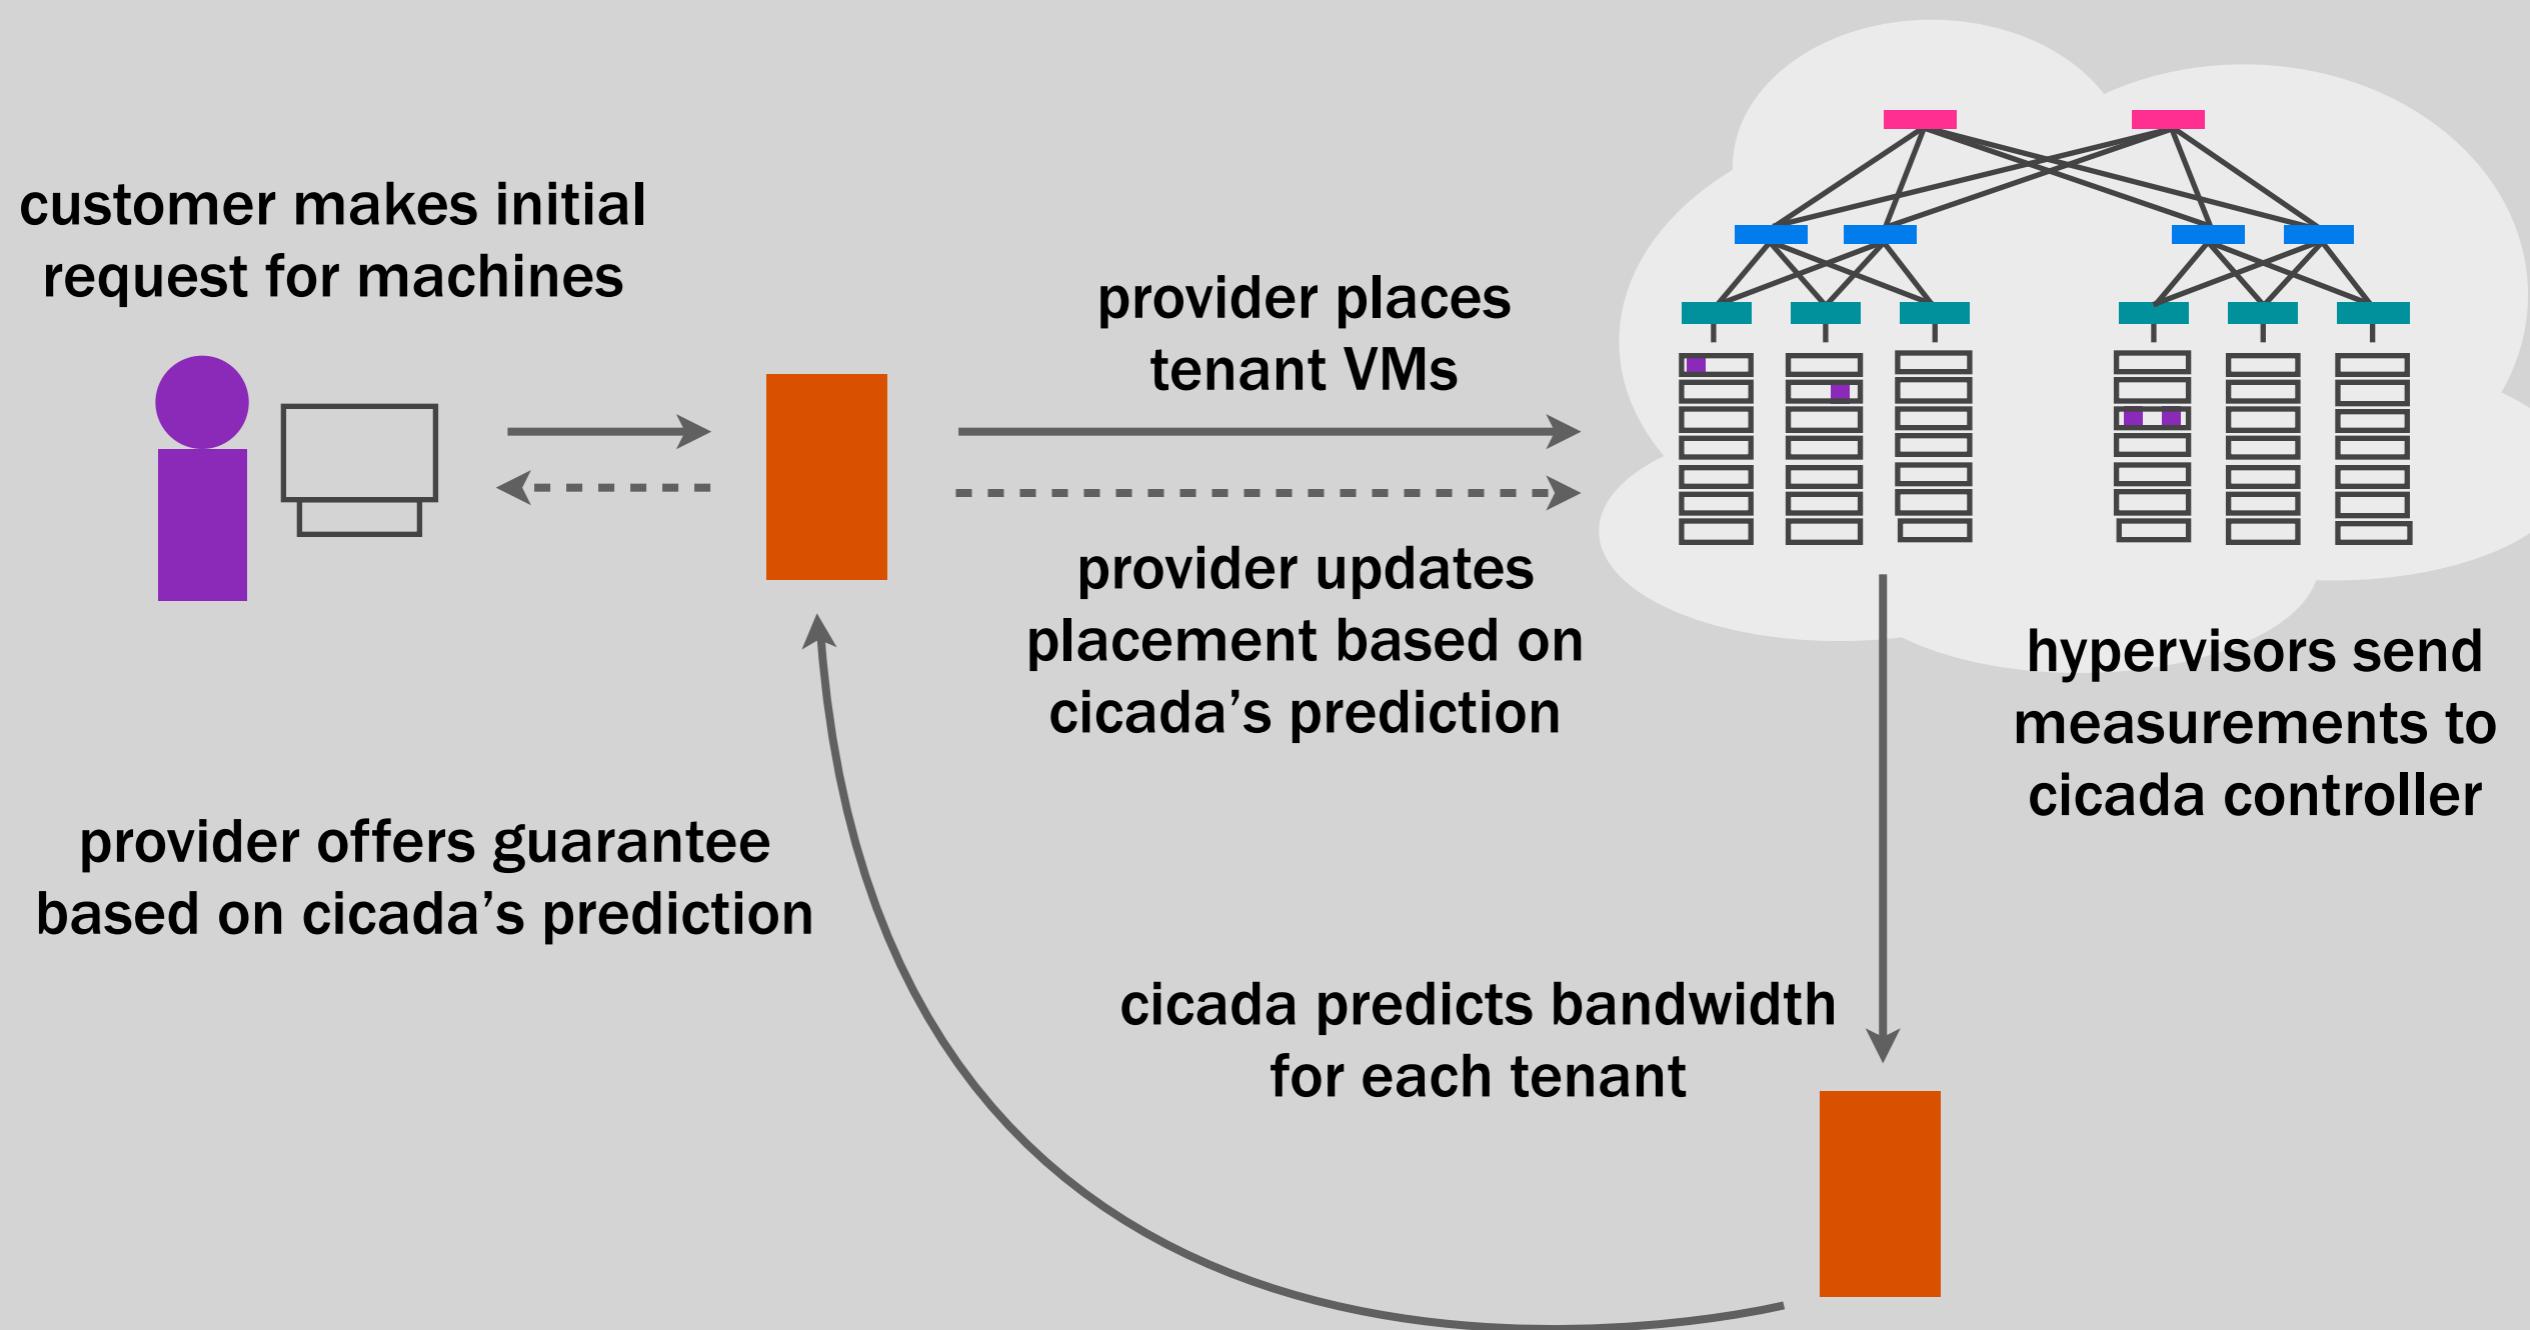
CICADA'S ARCHITECTURE

cicada makes **predictions** about an application's traffic to automatically generate a guarantee



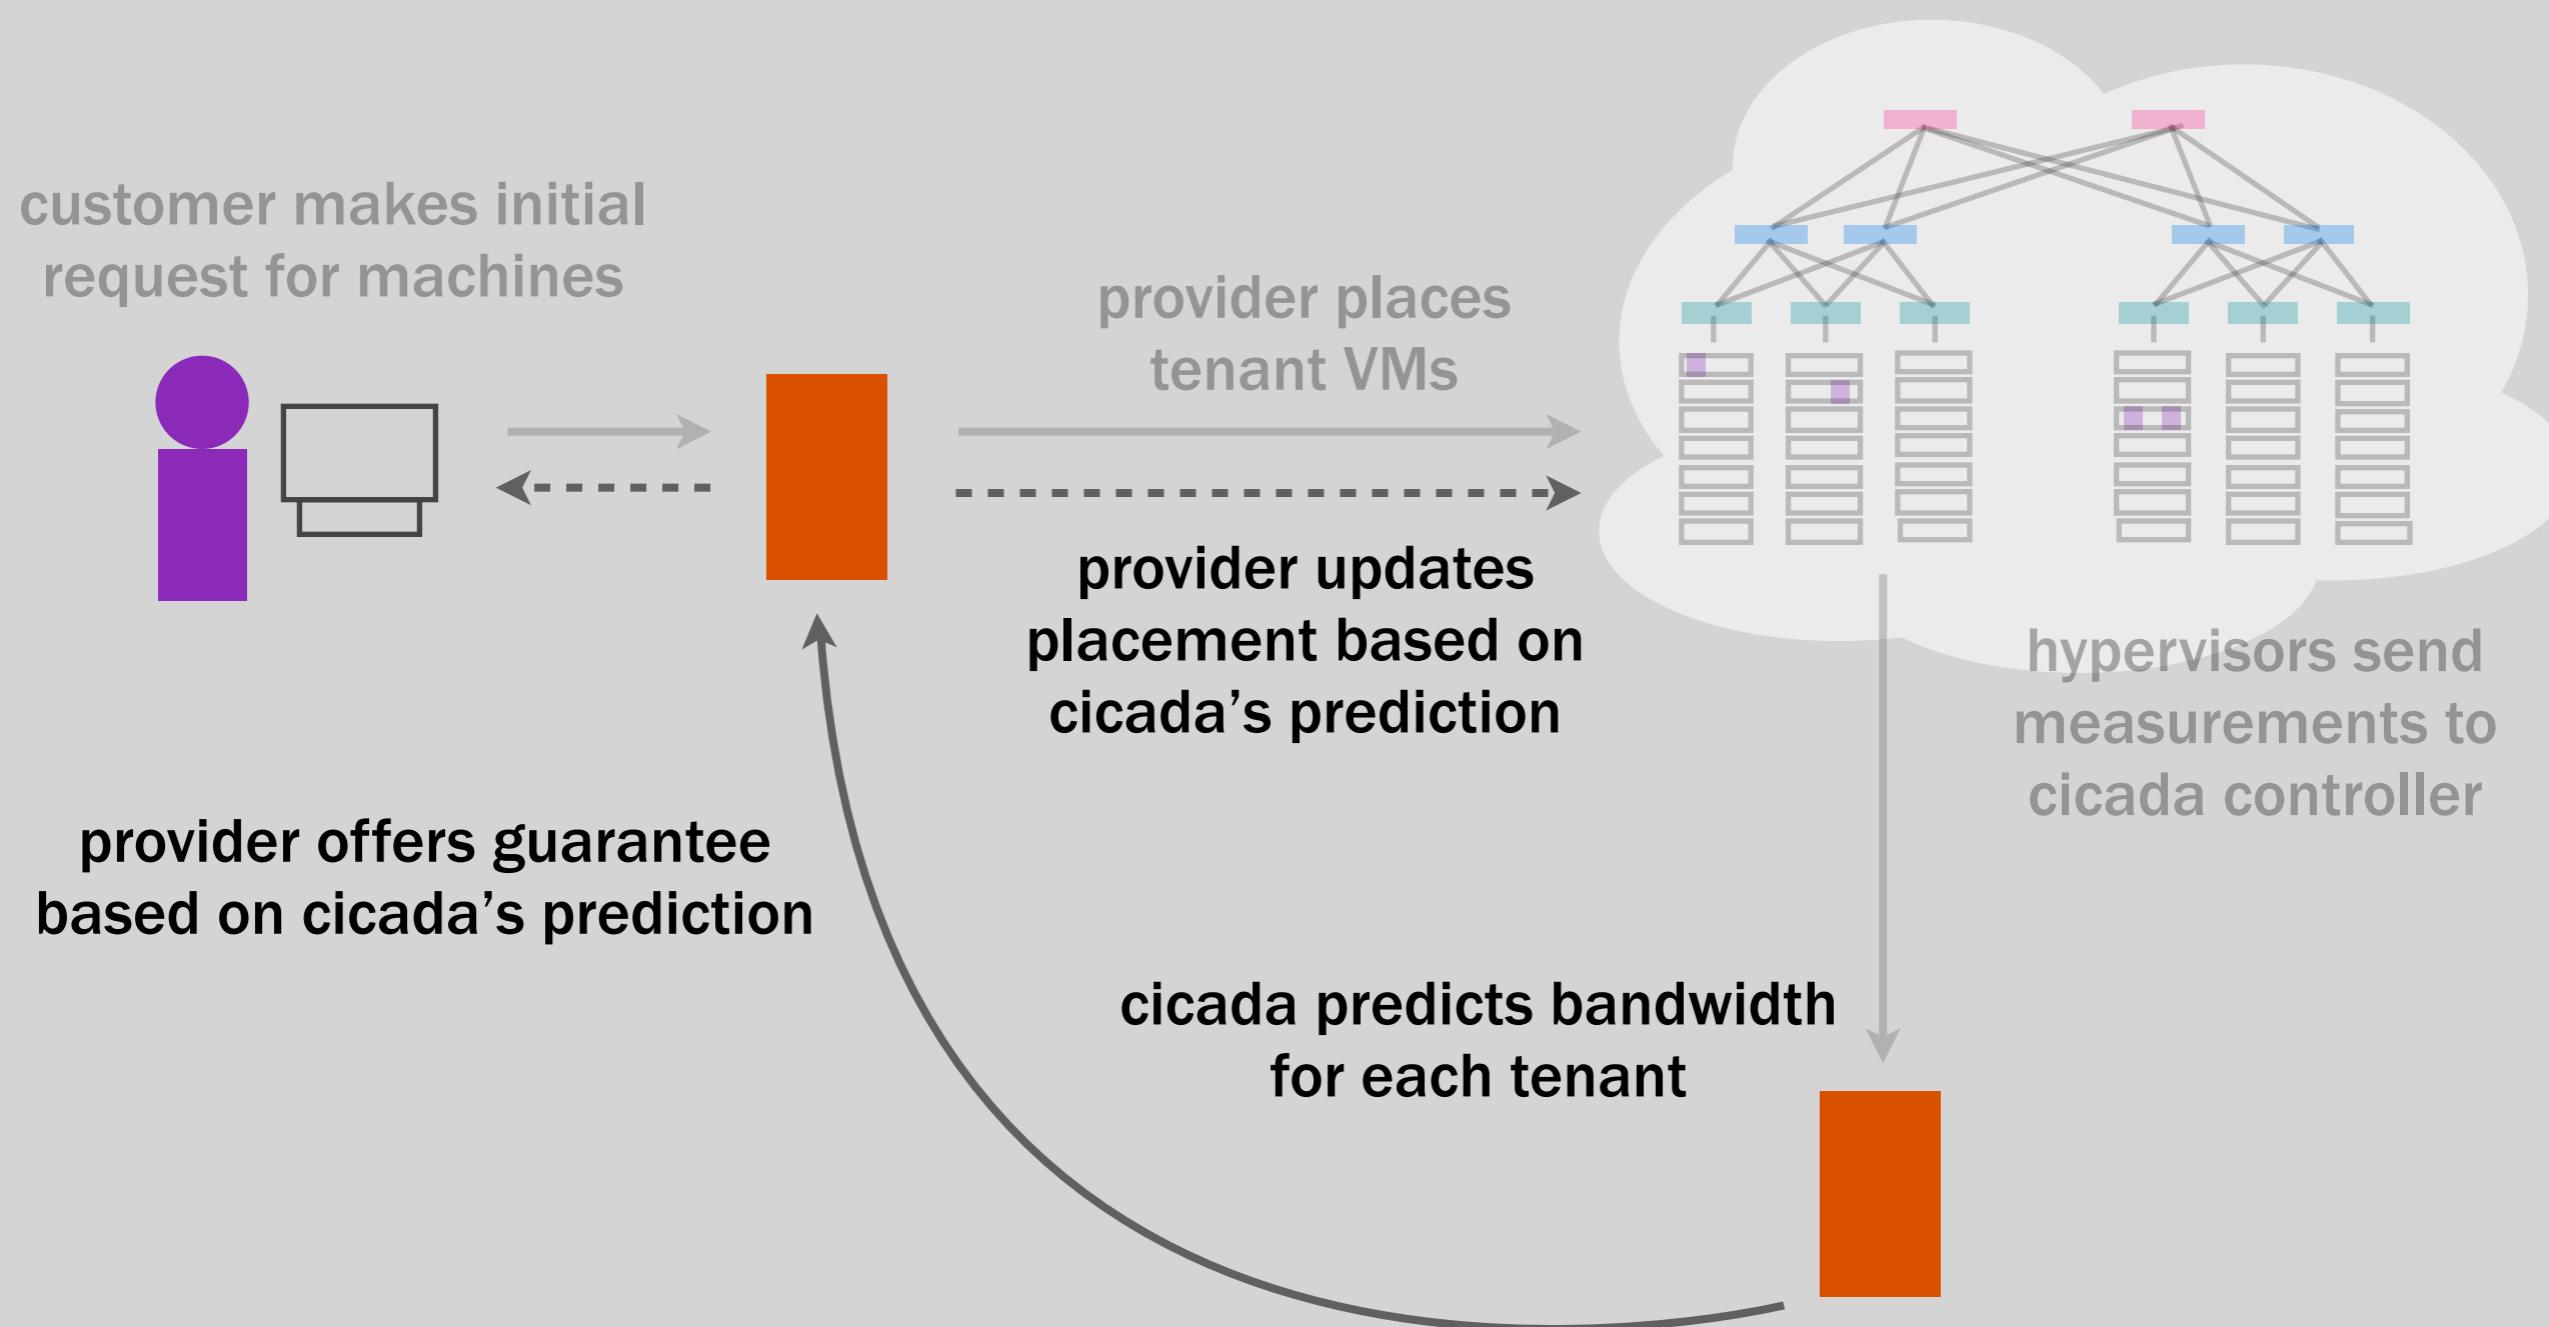
CICADA'S ARCHITECTURE

cicada makes **predictions** about an application's traffic to automatically generate a guarantee



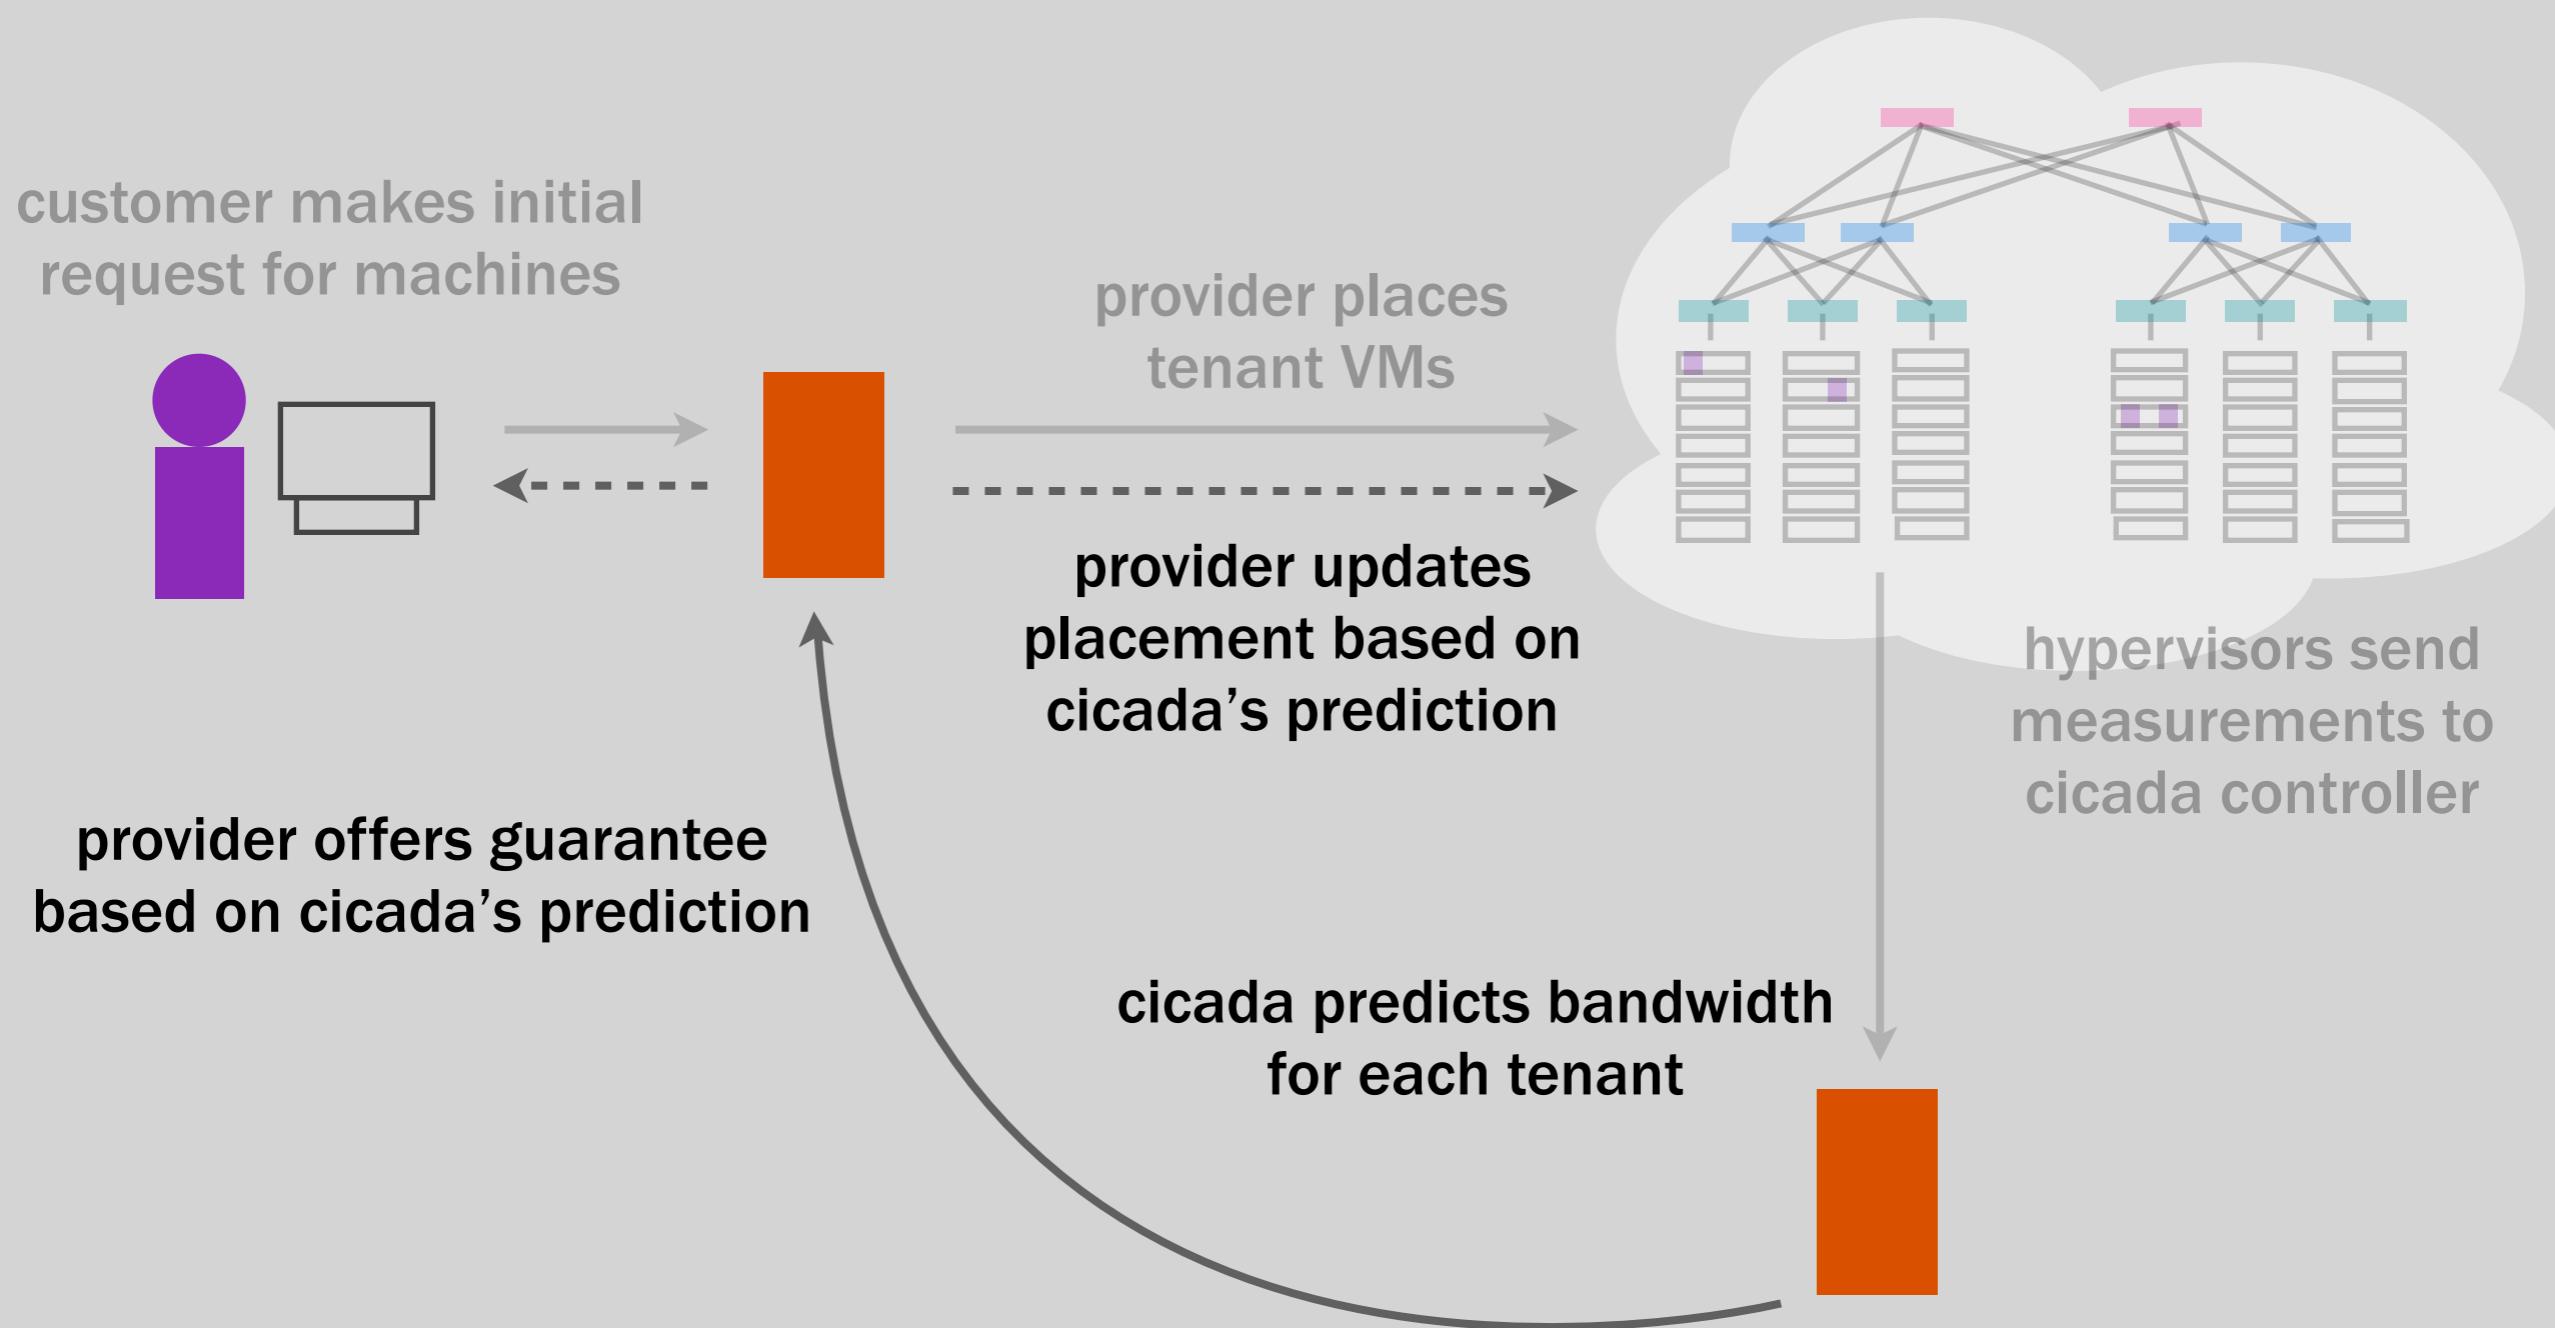
CICADA'S ARCHITECTURE

cicada makes **predictions** about an application's traffic to automatically generate a guarantee



CICADA'S ARCHITECTURE

problem: is application traffic actually predictable? if so, is it captured by existing models?



APPLICATION TRAFFIC

**to design cicada's prediction algorithm, we need to understand how
cloud applications send traffic**

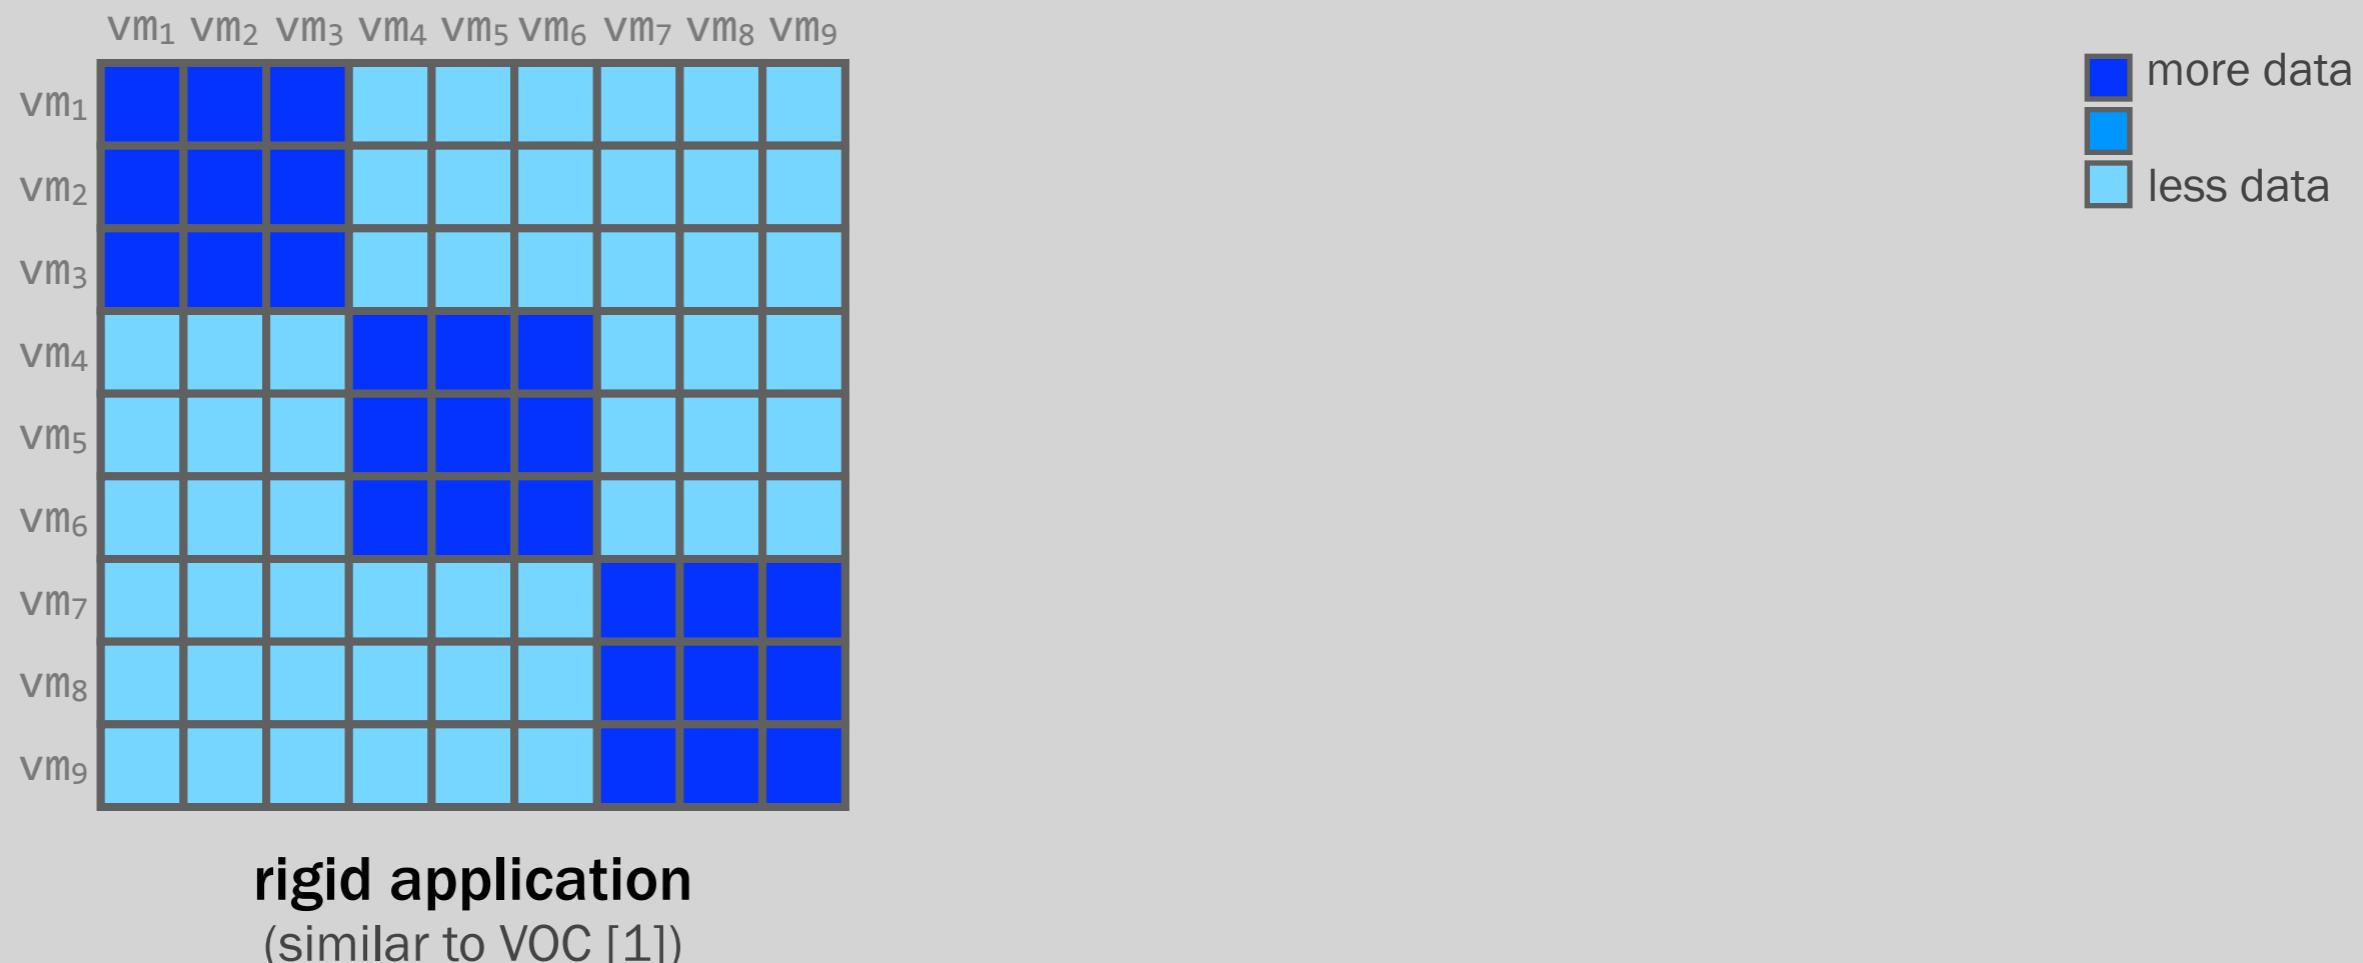
APPLICATION TRAFFIC

to design cicada's prediction algorithm, we need to understand how cloud applications send traffic



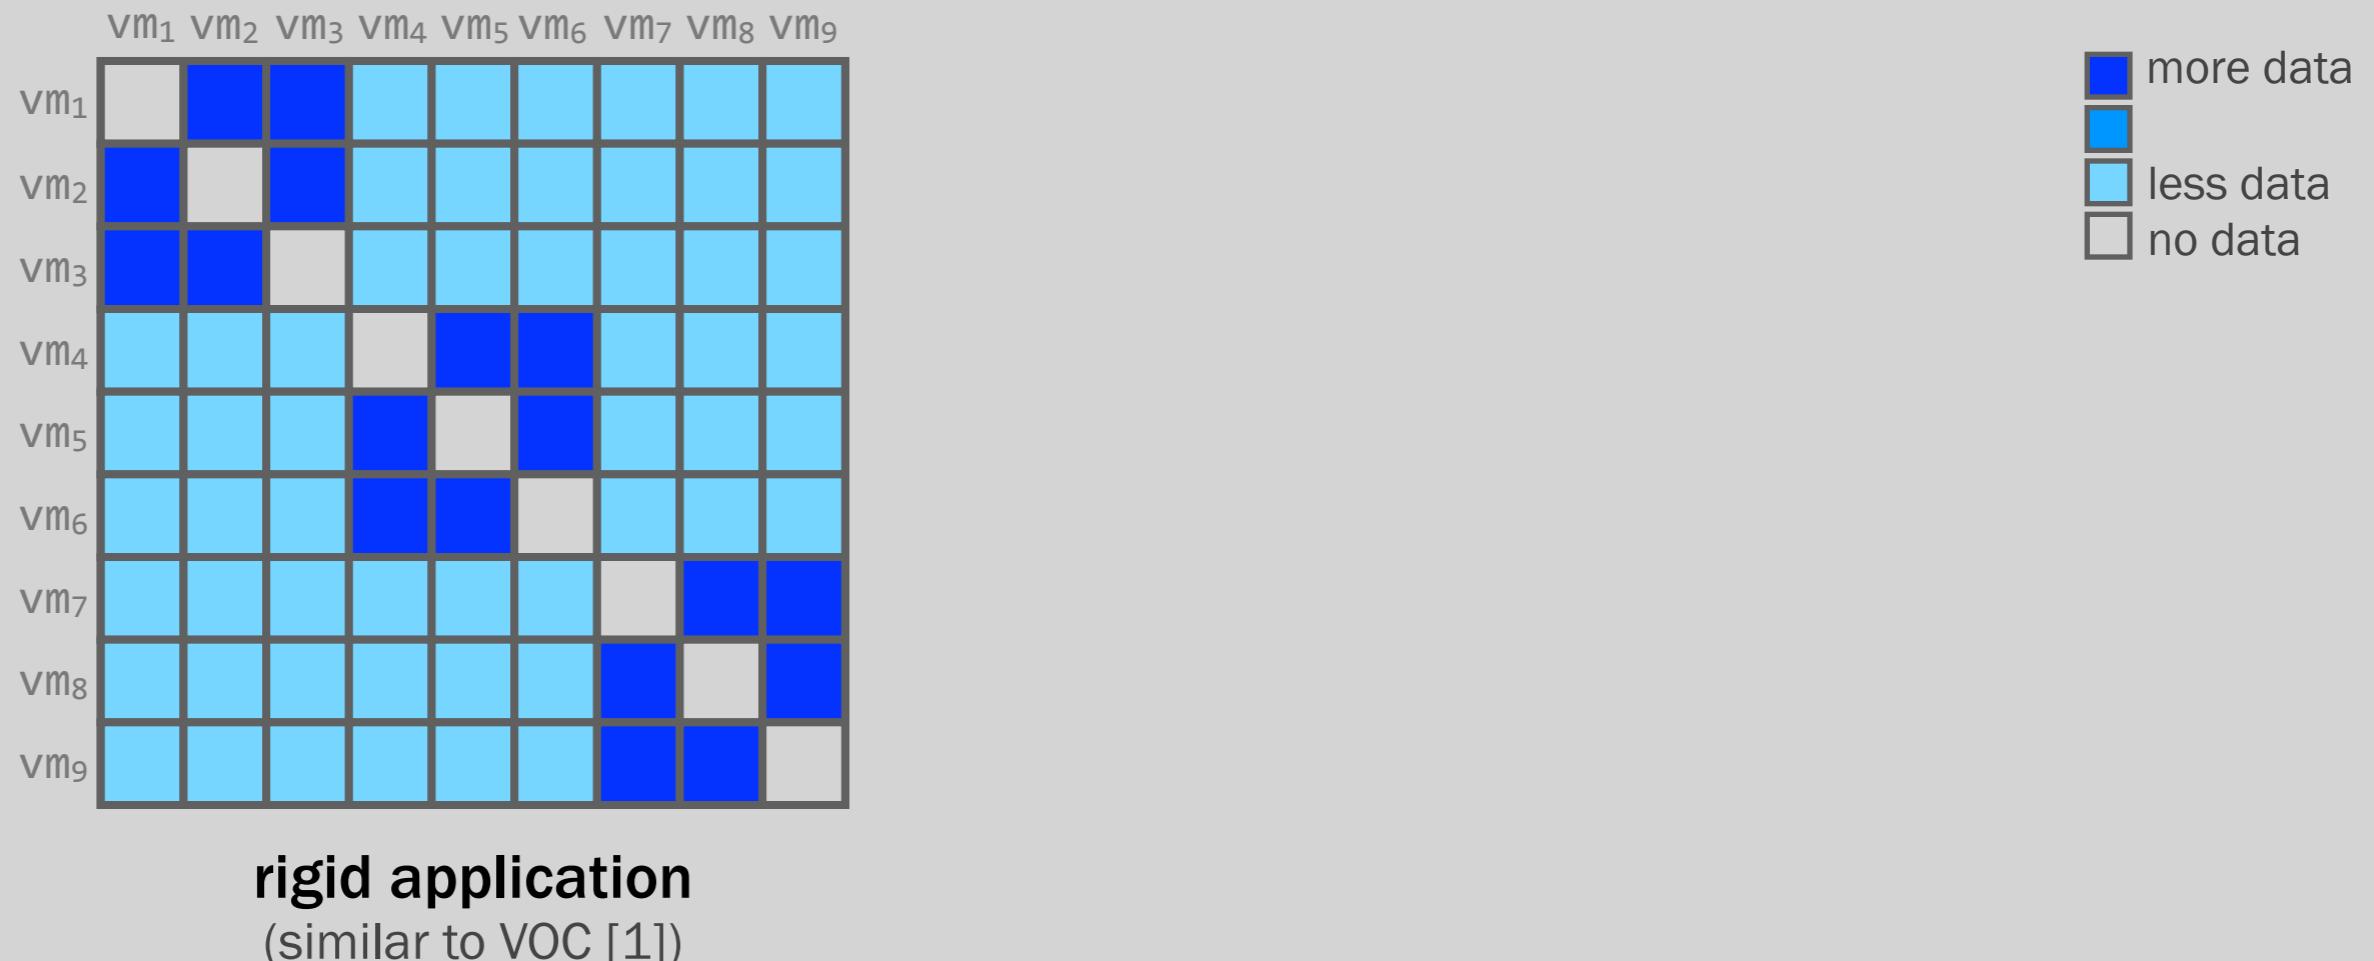
APPLICATION TRAFFIC

to design cicada's prediction algorithm, we need to understand how cloud applications send traffic



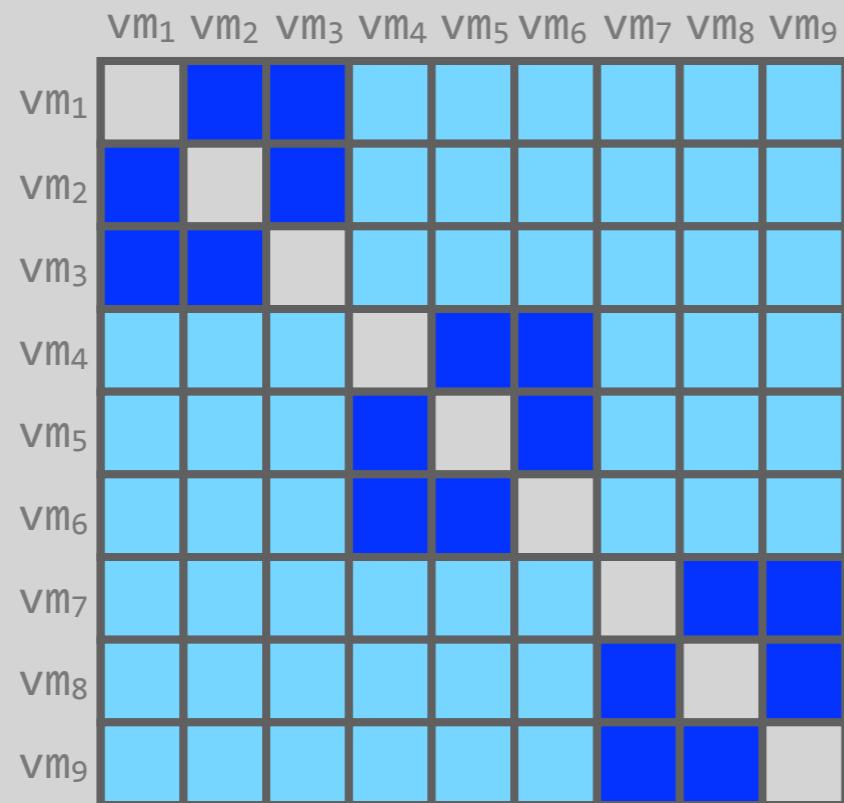
APPLICATION TRAFFIC

to design cicada's prediction algorithm, we need to understand how cloud applications send traffic

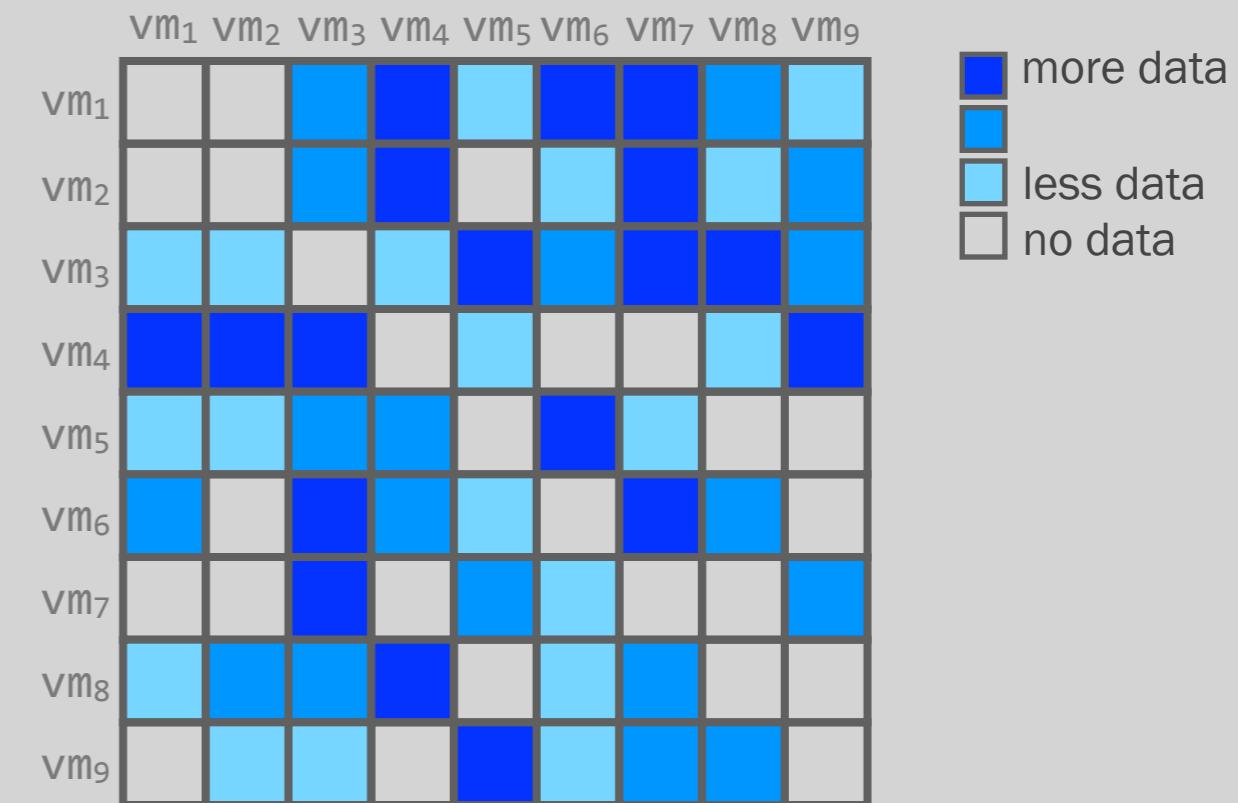


APPLICATION TRAFFIC

to design cicada's prediction algorithm, we need to understand how cloud applications send traffic



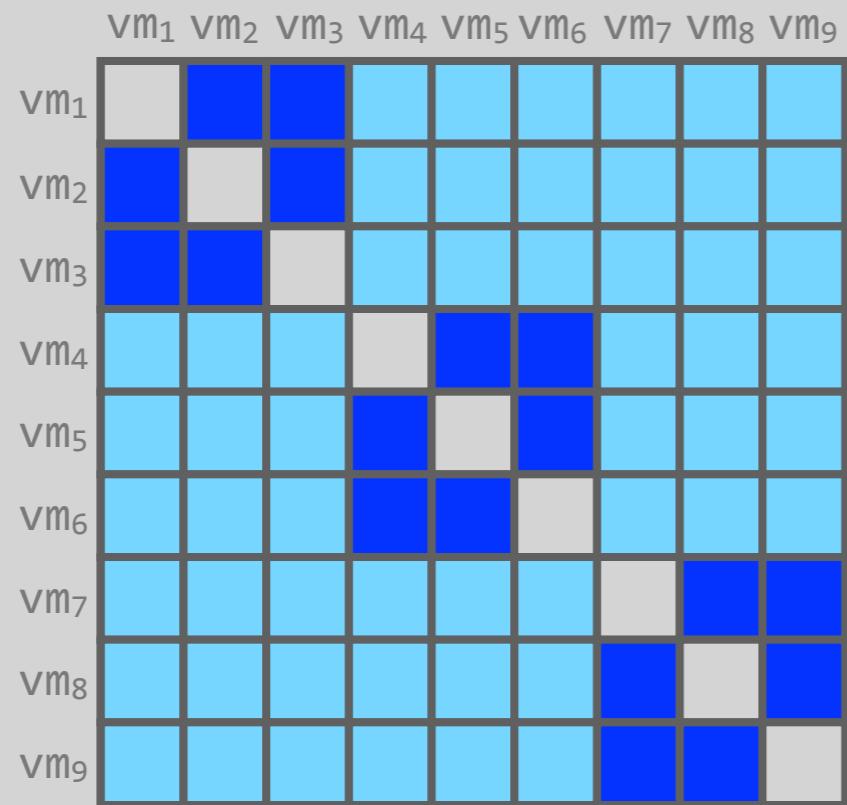
rigid application
(similar to VOC [1])



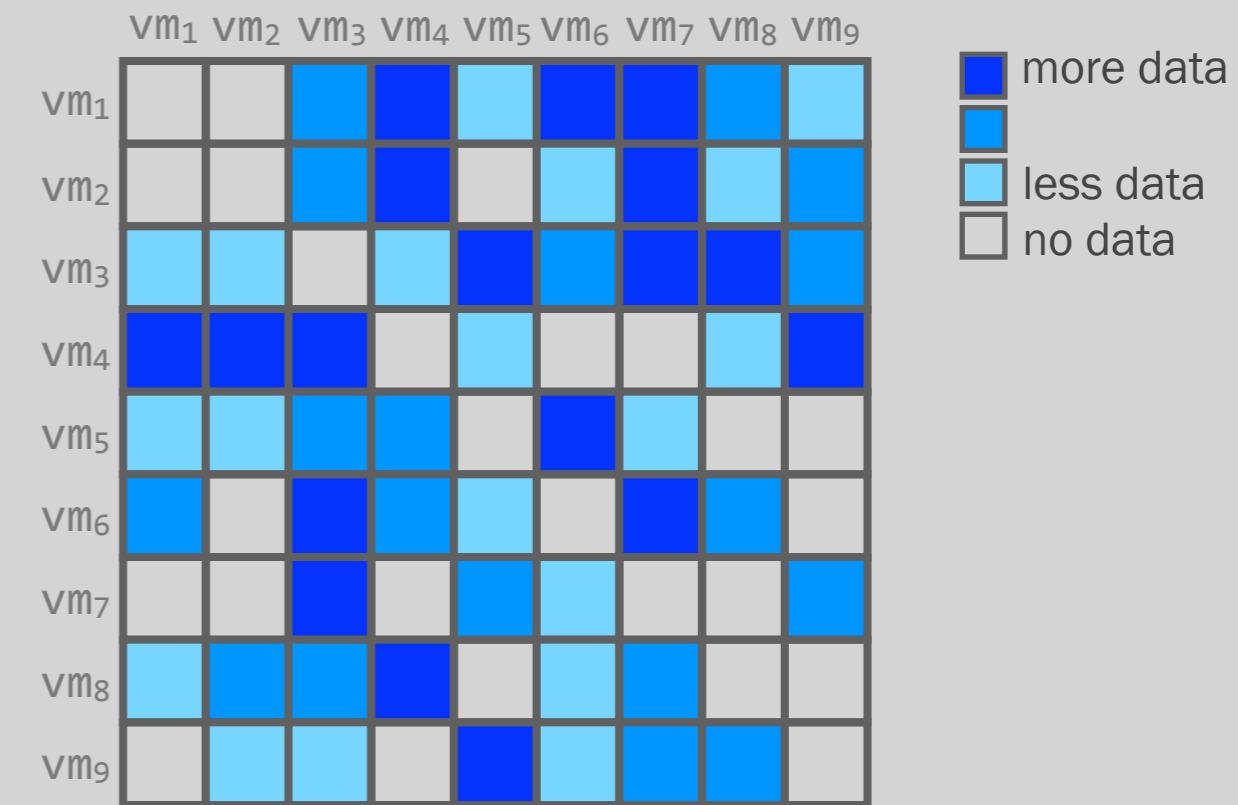
variable application

APPLICATION TRAFFIC

to design cicada's prediction algorithm, we need to understand how cloud applications send traffic



rigid application
(similar to VOC [1])

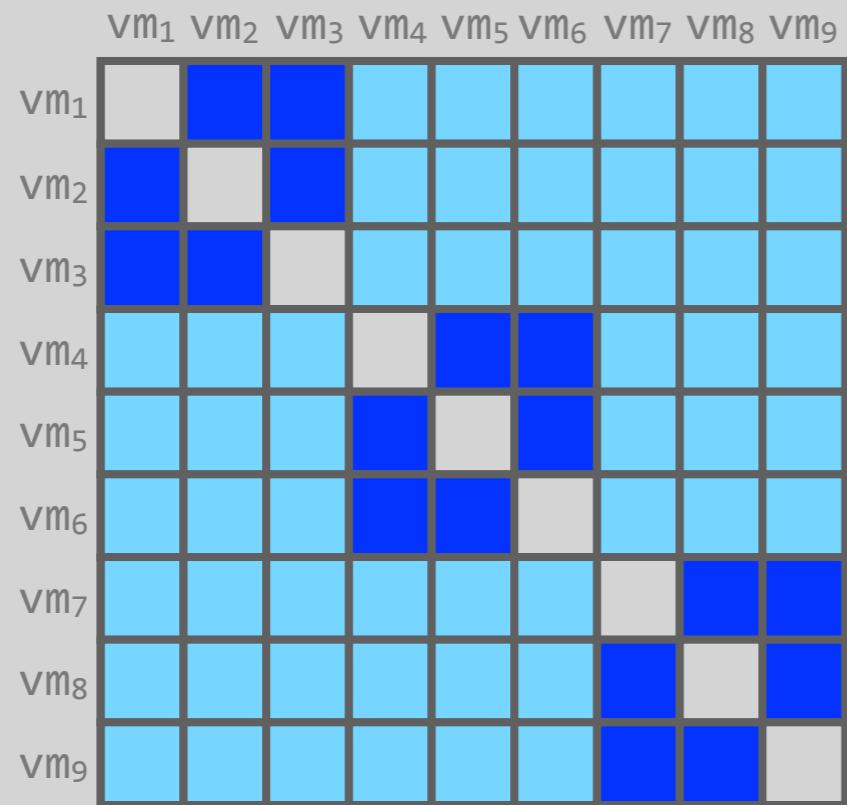


variable application

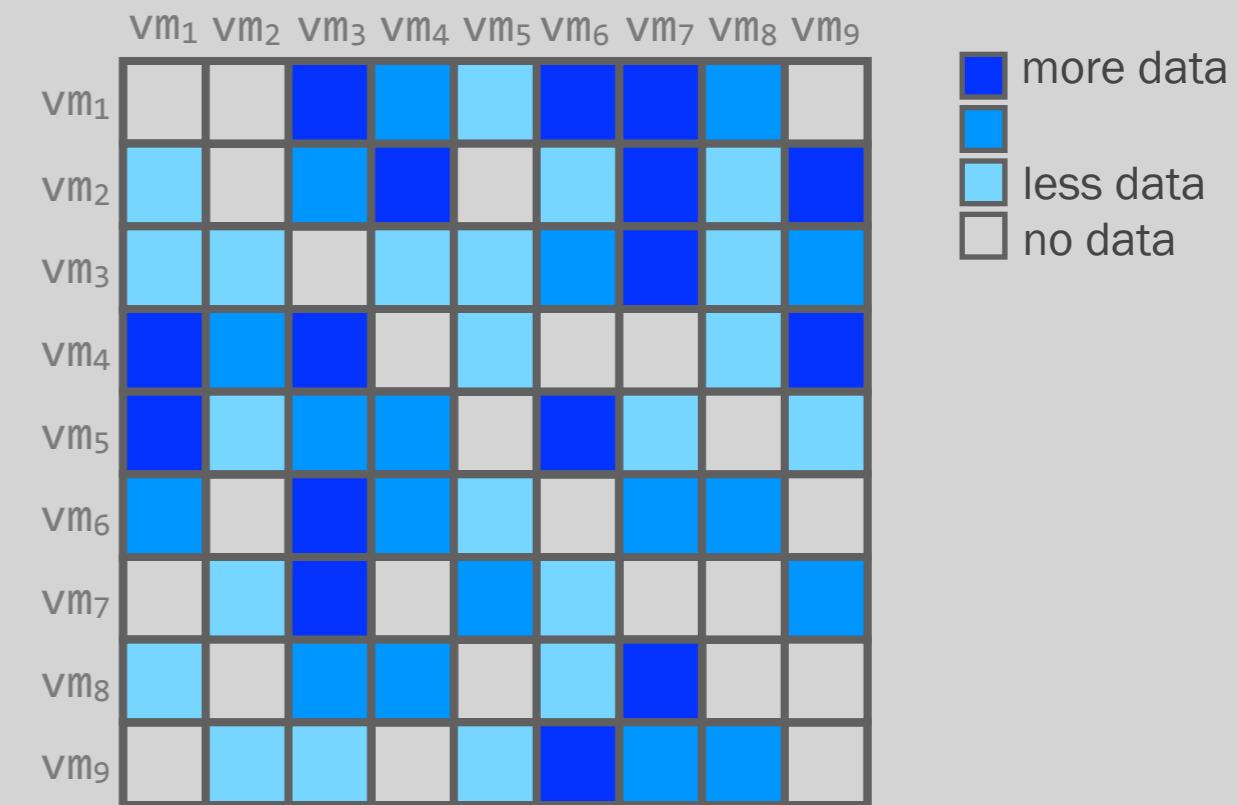
spatial variability: different pairs of VMs transfer different amounts of data

APPLICATION TRAFFIC

to design cicada's prediction algorithm, we need to understand how cloud applications send traffic



rigid application
(similar to VOC [1])

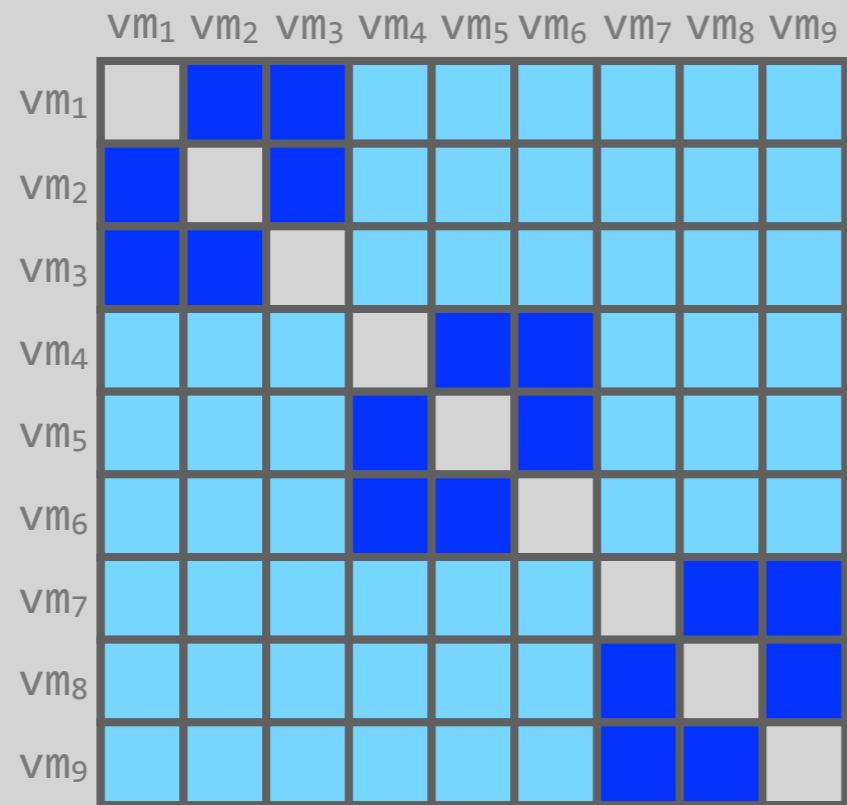


variable application

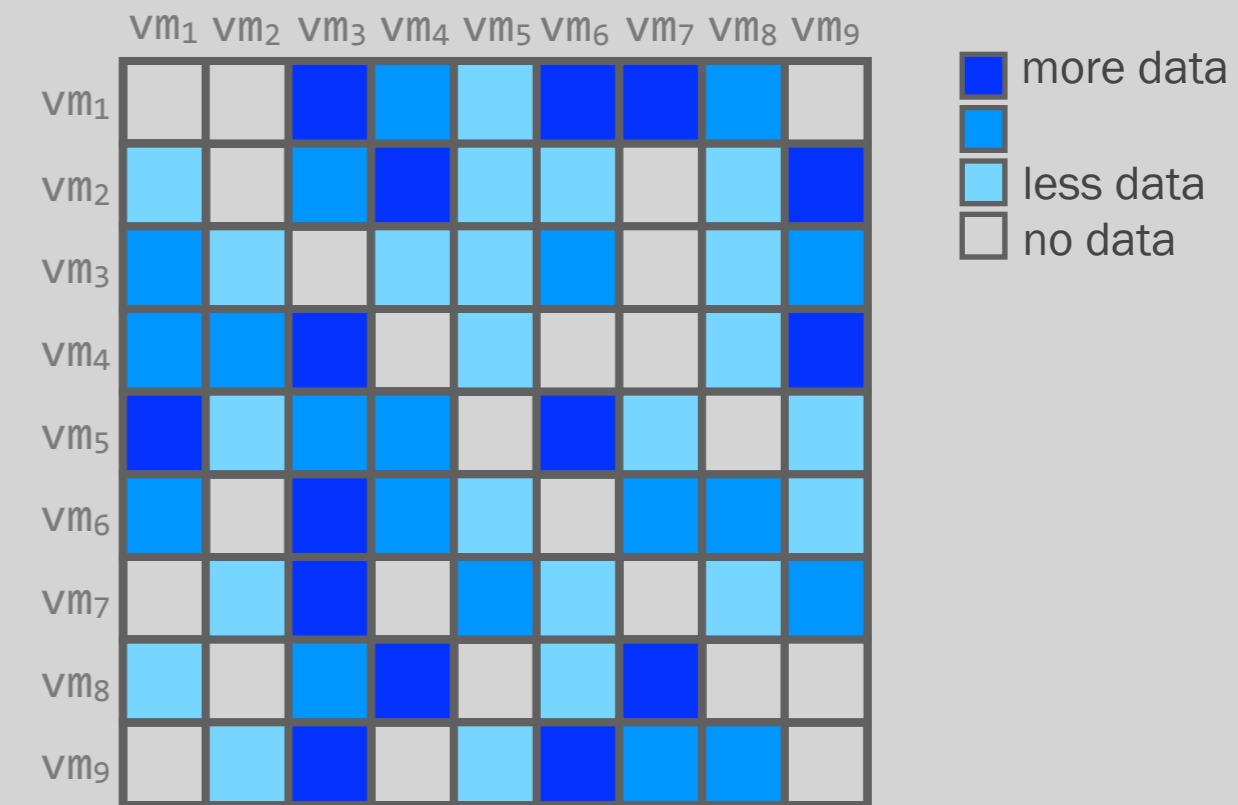
spatial variability: different pairs of VMs transfer different amounts of data

APPLICATION TRAFFIC

to design cicada's prediction algorithm, we need to understand how cloud applications send traffic



rigid application
(similar to VOC [1])

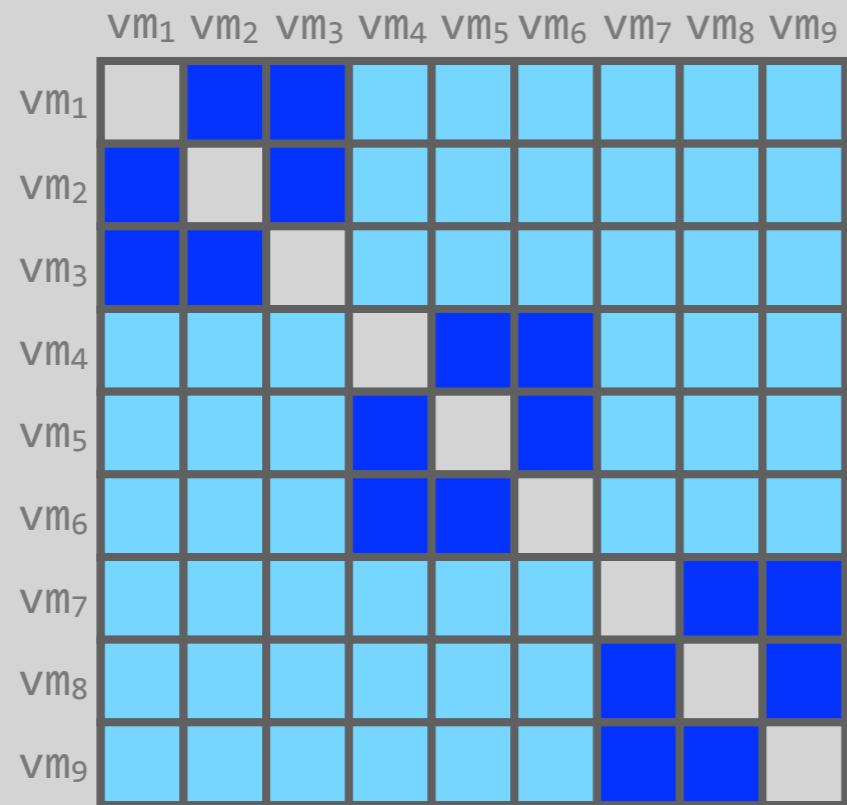


variable application

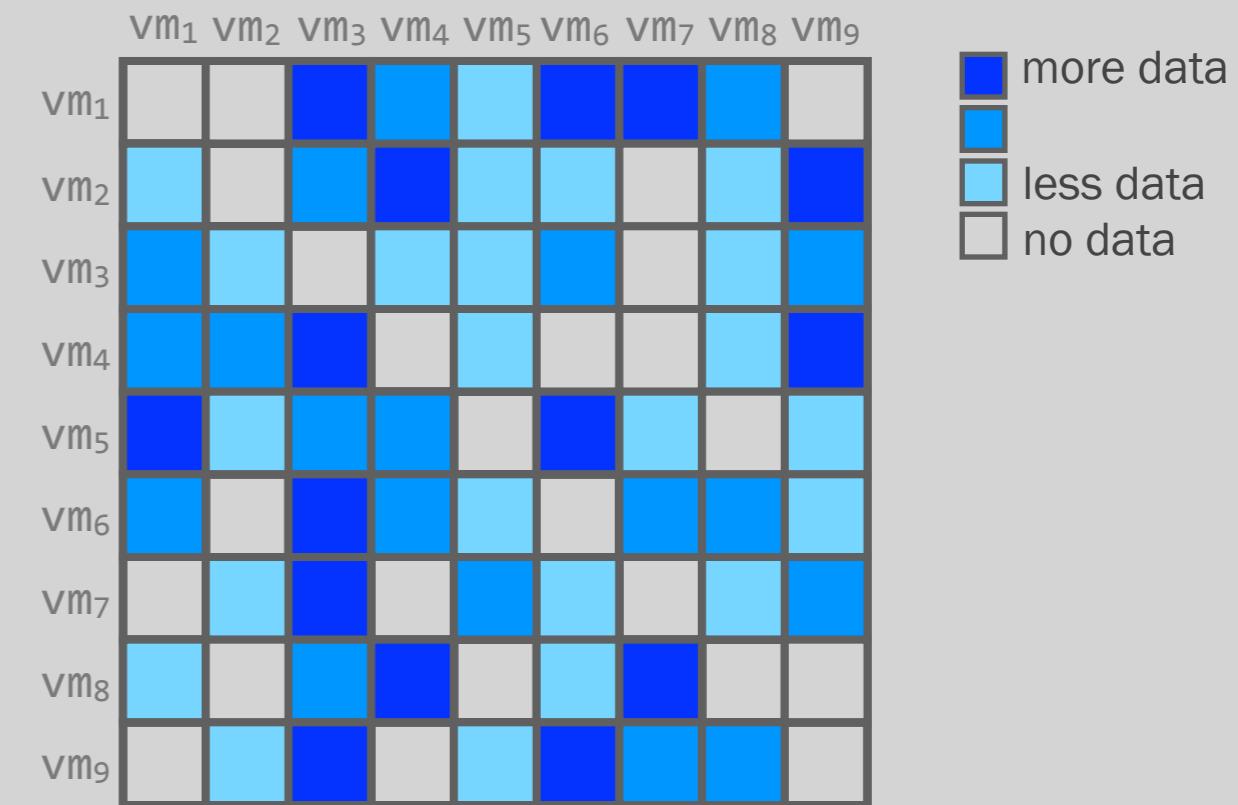
spatial variability: different pairs of VMs transfer different amounts of data

APPLICATION TRAFFIC

to design cicada's prediction algorithm, we need to understand how cloud applications send traffic



rigid application
(similar to VOC [1])

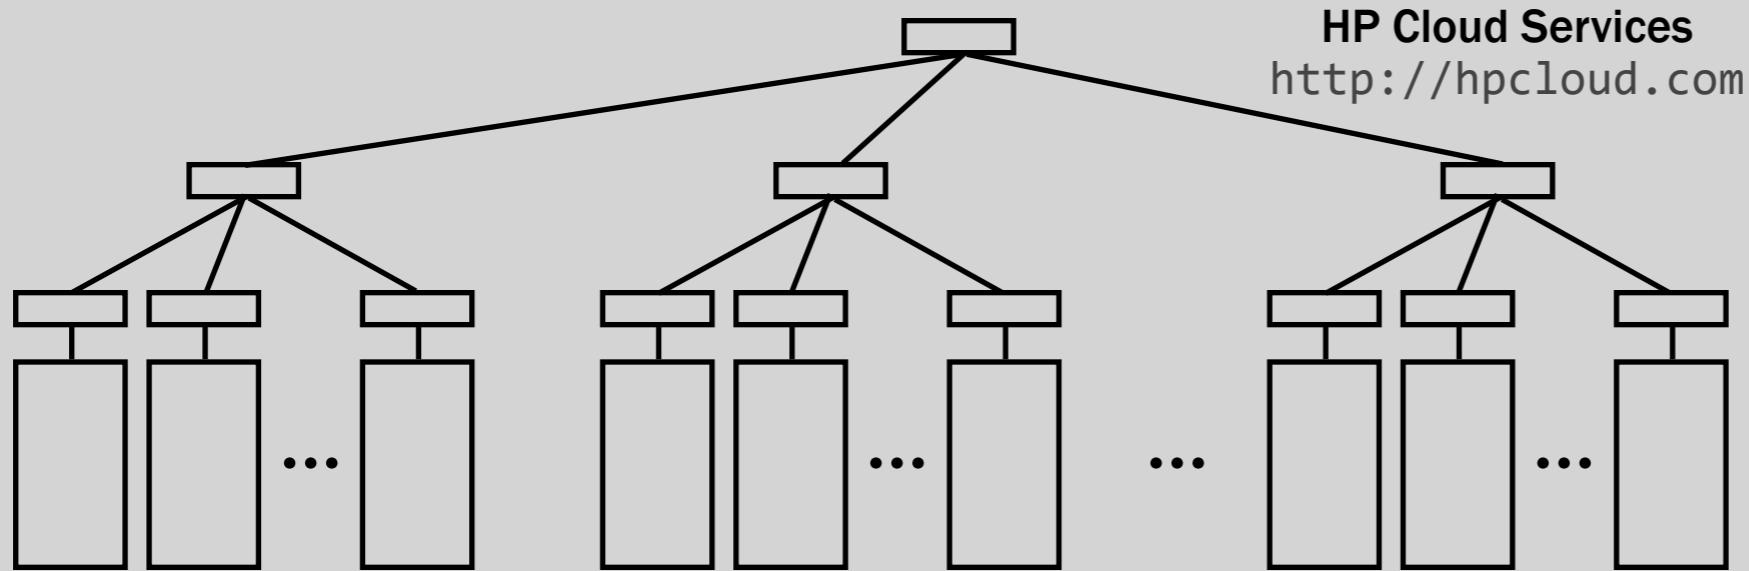


variable application

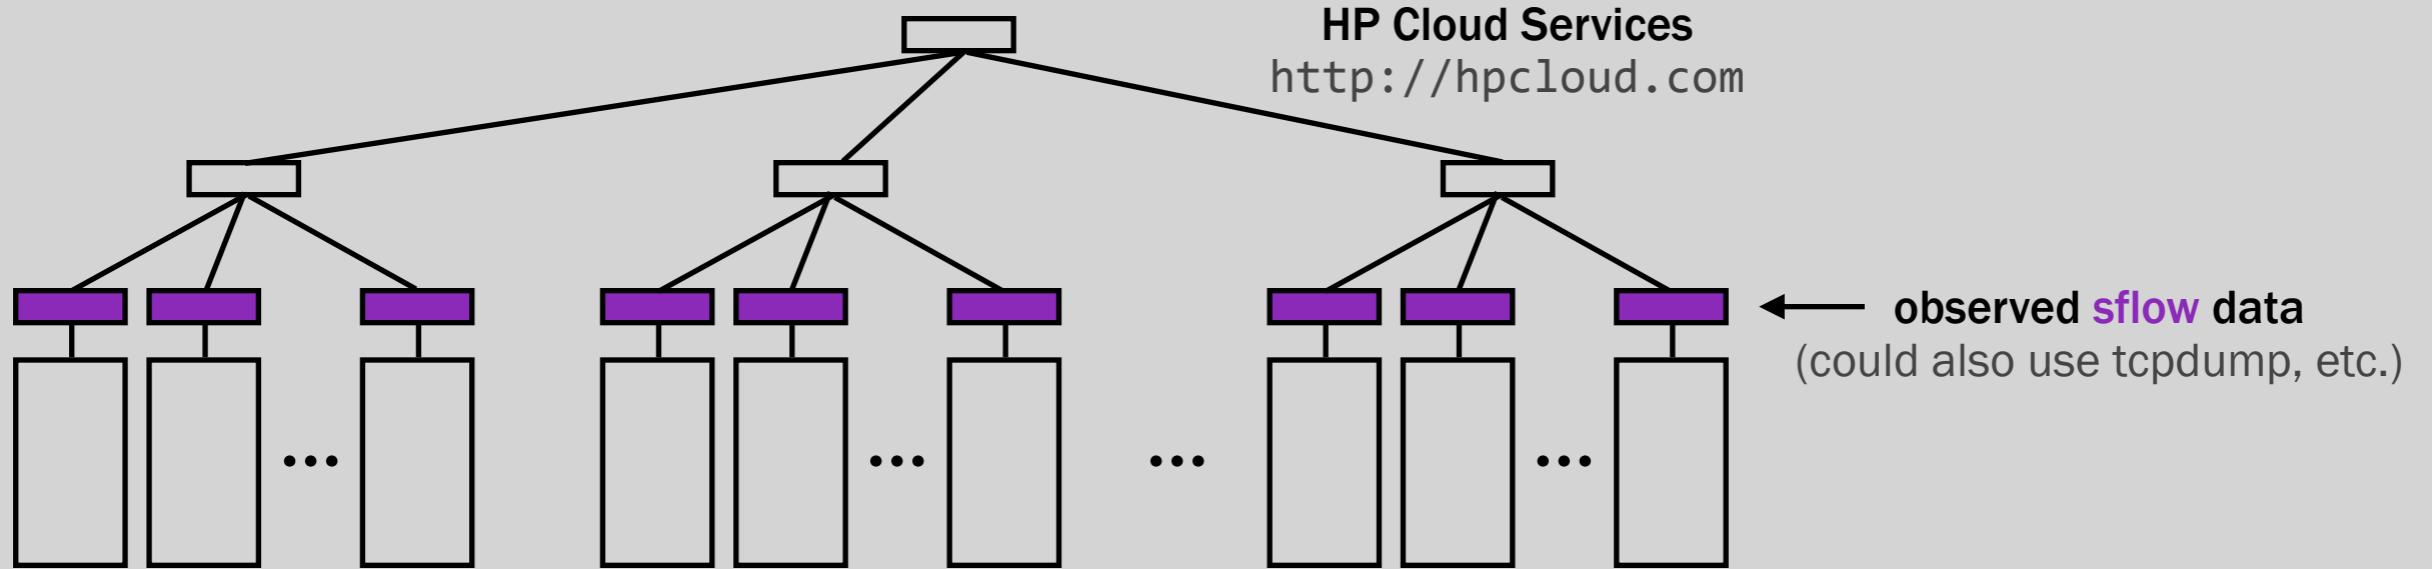
spatial variability: different pairs of VMs transfer different amounts of data

temporal variability: pairs of VMs transfer different amounts of data at different times

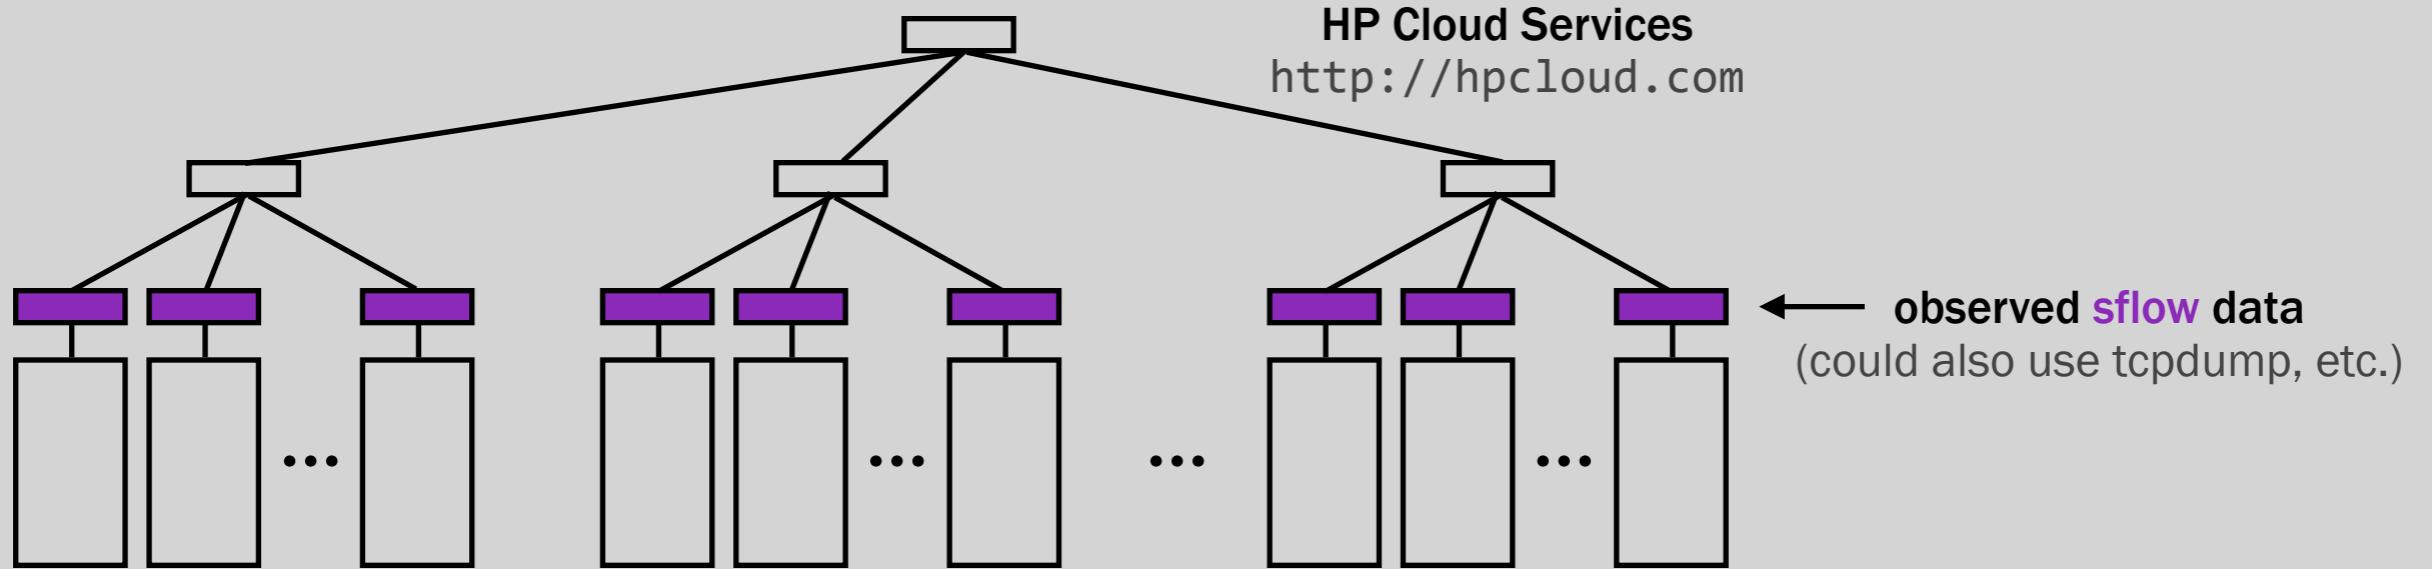
DATASET



DATASET



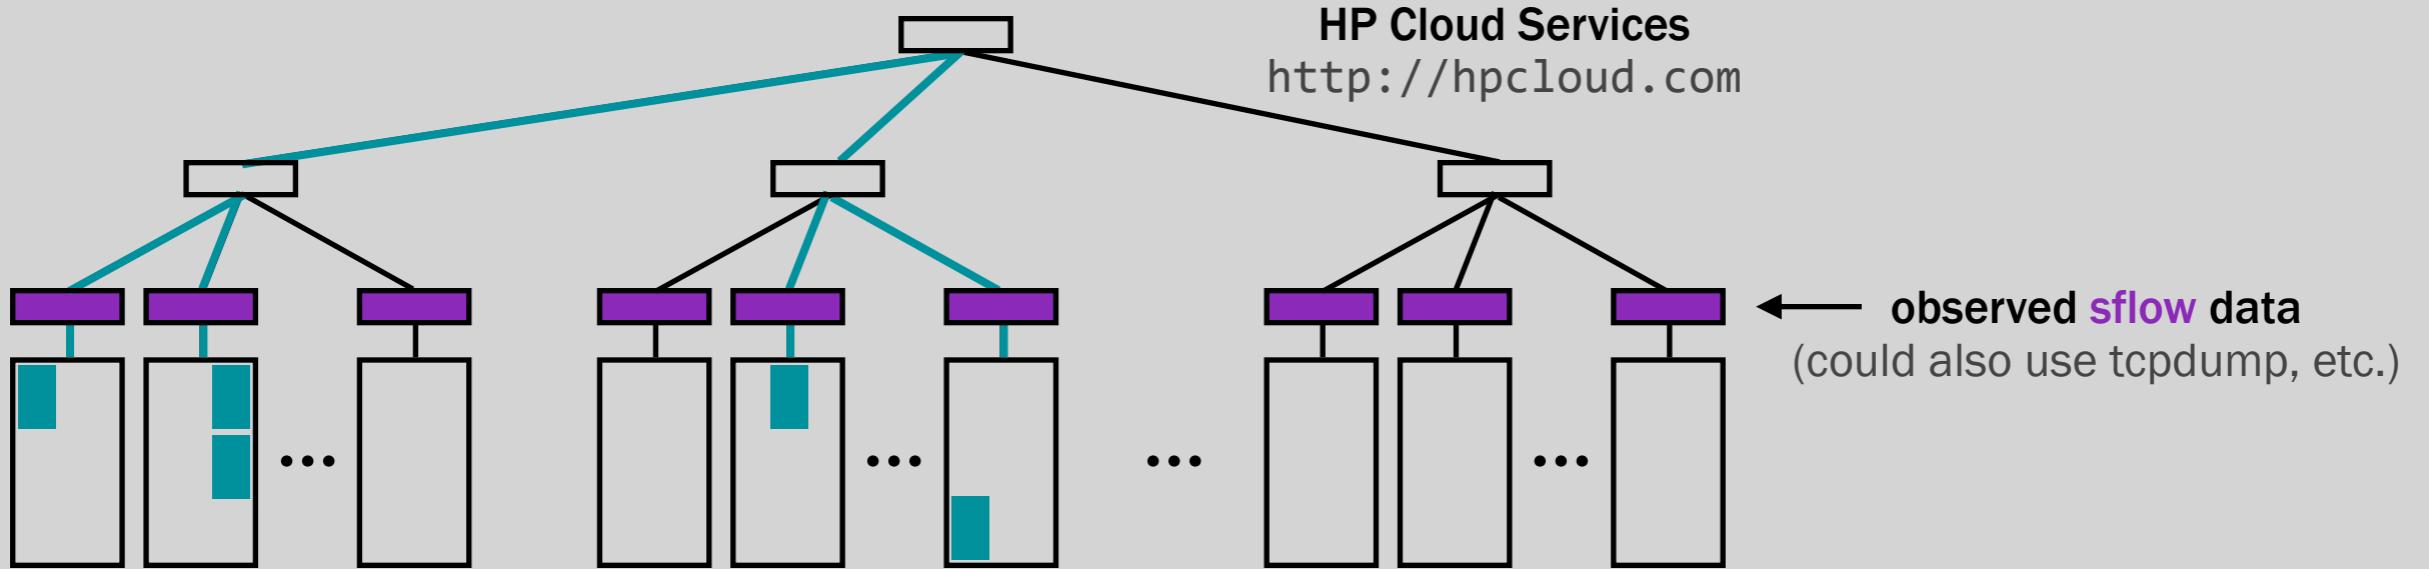
DATASET



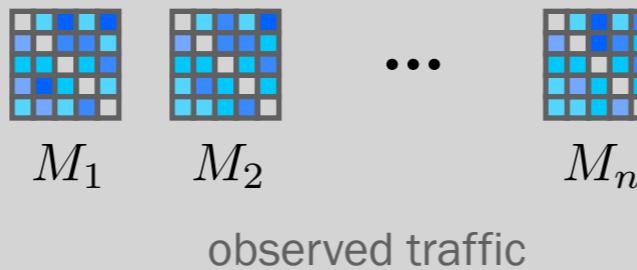
(in this talk)

**collected one traffic
matrix per hour, for
each application**

DATASET

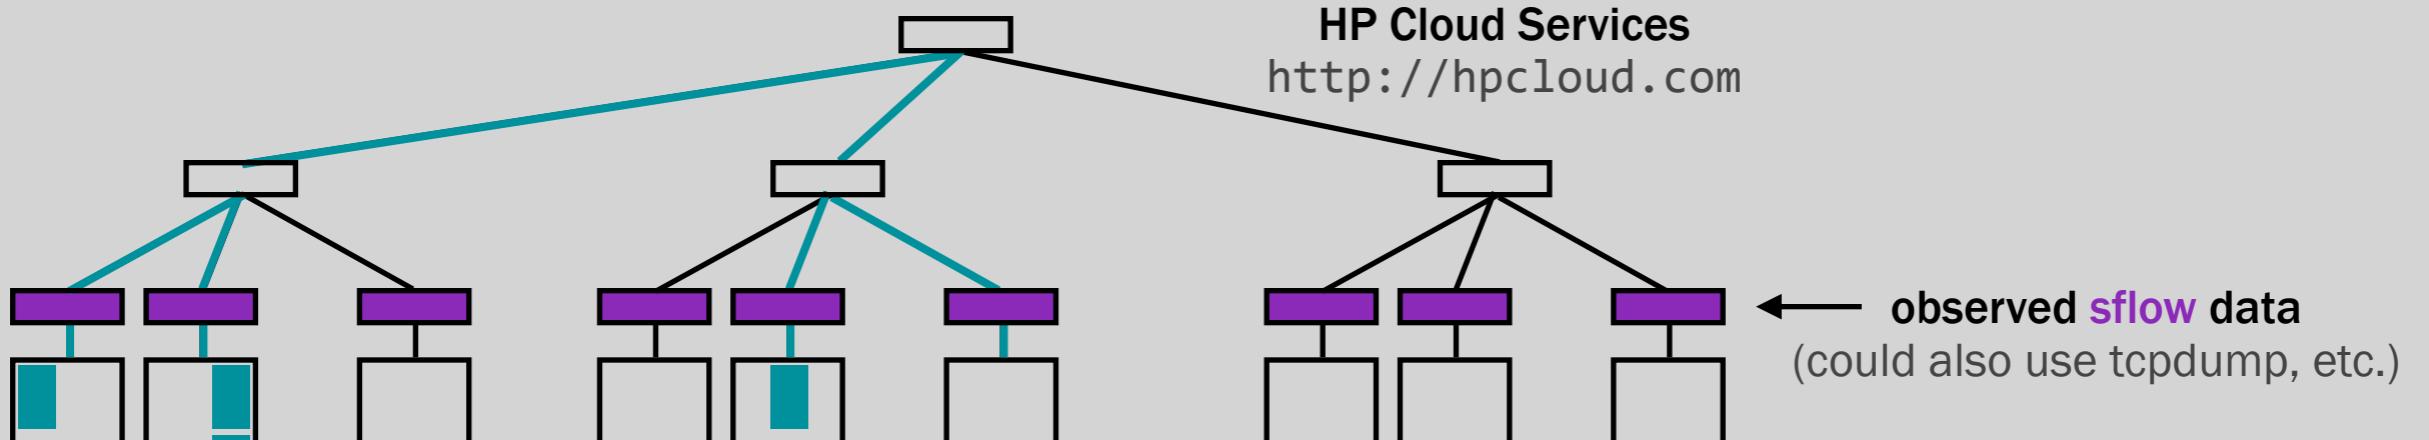


(in this talk)
collected one traffic matrix per hour, for each application

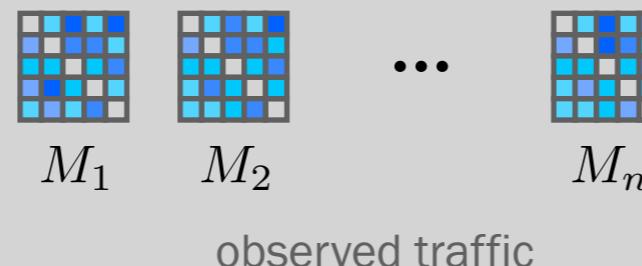


each entry represents the **number of bytes** transferred between i and j

DATASET



(in this talk)
collected one traffic matrix per hour, for each application



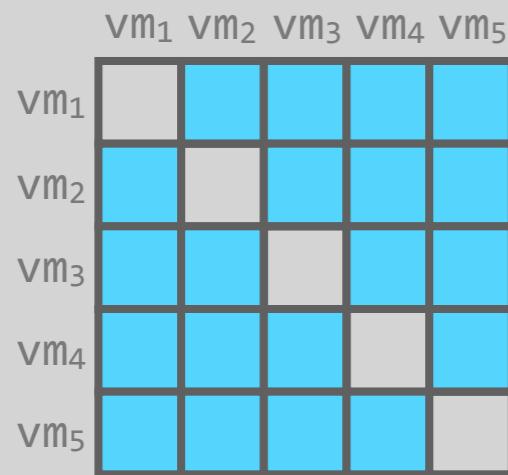
each entry represents the number of bytes transferred between i and j

goal: use this dataset to quantify spatial and temporal variability

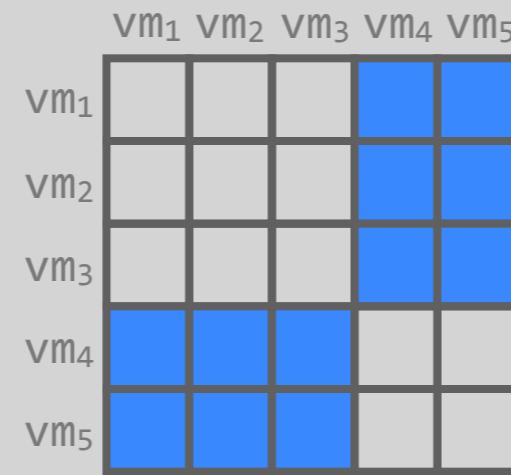
SPATIAL VARIABILITY

how do we quantify spatial variability?

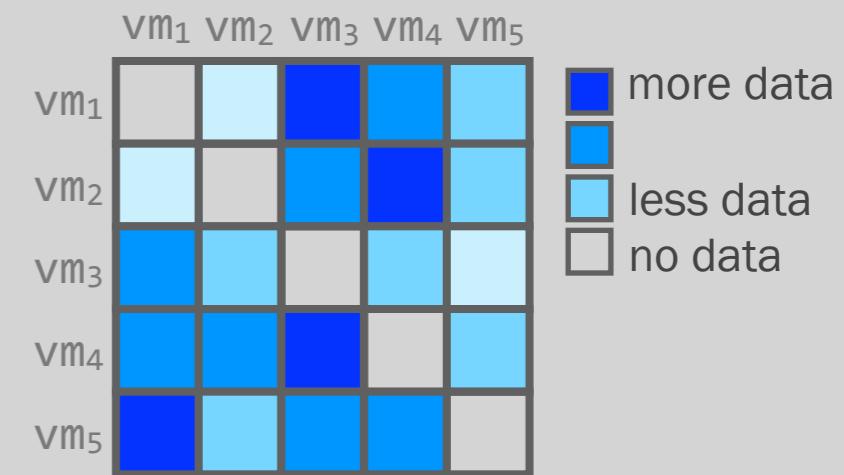
no spatial variability



some spatial variability



high spatial variability

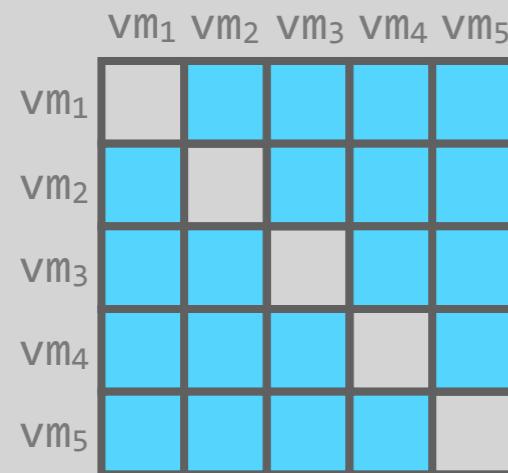


- █ more data
- █ less data
- █ no data

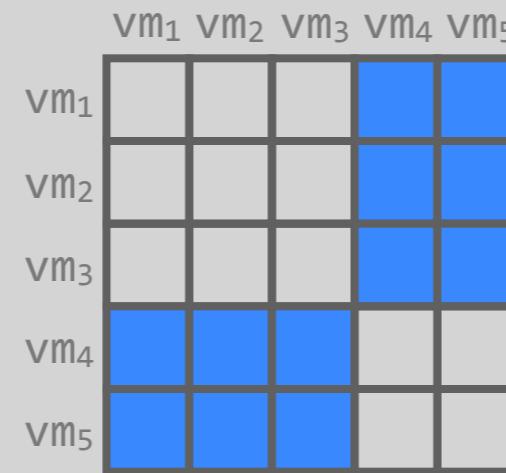
SPATIAL VARIABILITY

how do we quantify spatial variability?

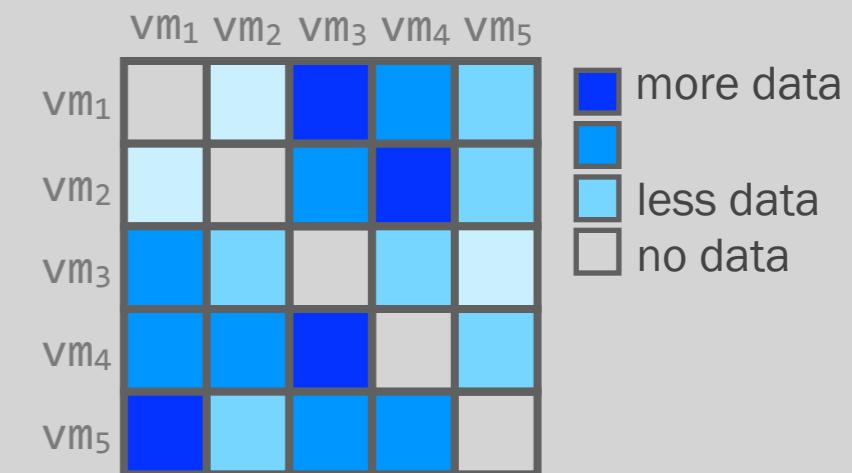
no spatial variability



some spatial variability



high spatial variability

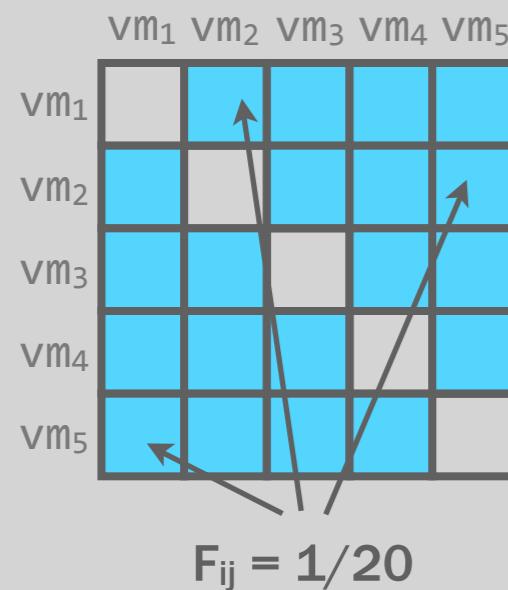


1. let F_{ij} = fraction of tenant traffic sent from VM_i to VM_j

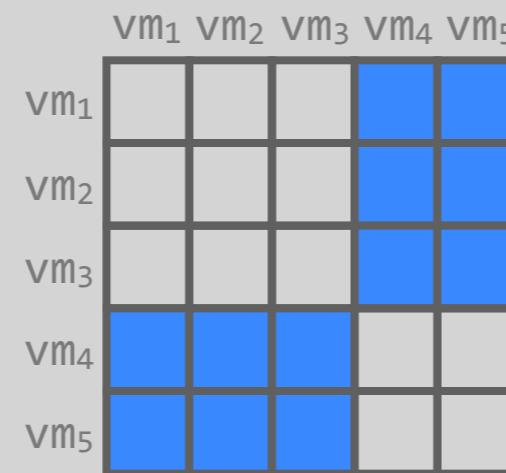
SPATIAL VARIABILITY

how do we quantify spatial variability?

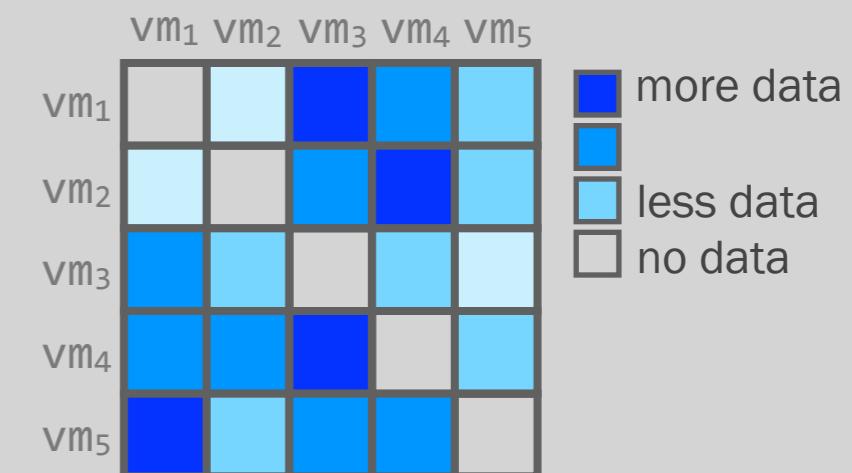
no spatial variability



some spatial variability



high spatial variability

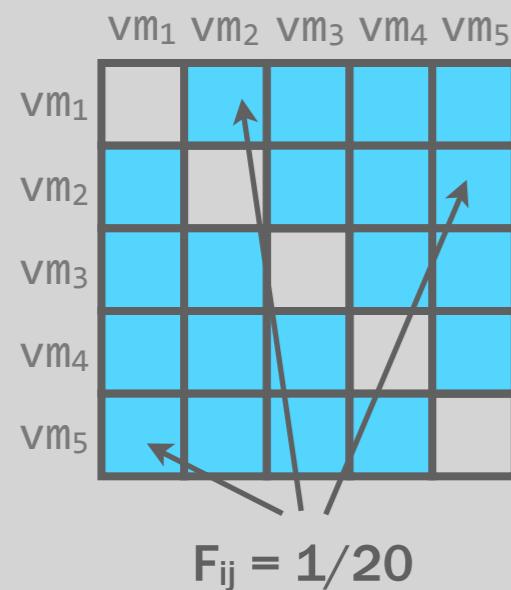


1. let F_{ij} = fraction of tenant traffic sent from VM_i to VM_j

SPATIAL VARIABILITY

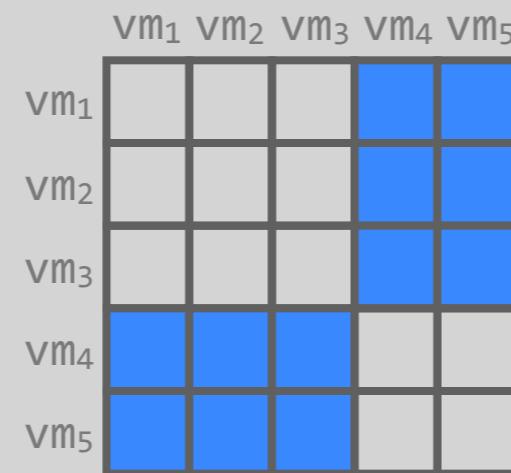
how do we quantify spatial variability?

no spatial variability

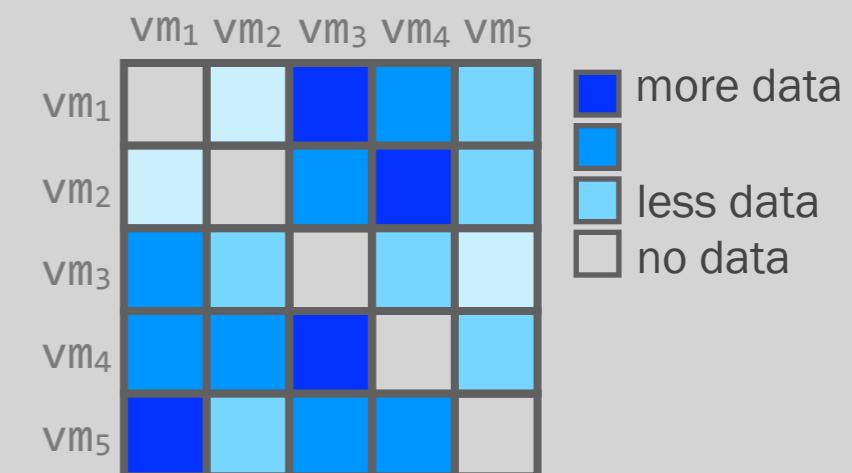


all F_{ij} values are equal

some spatial variability



high spatial variability

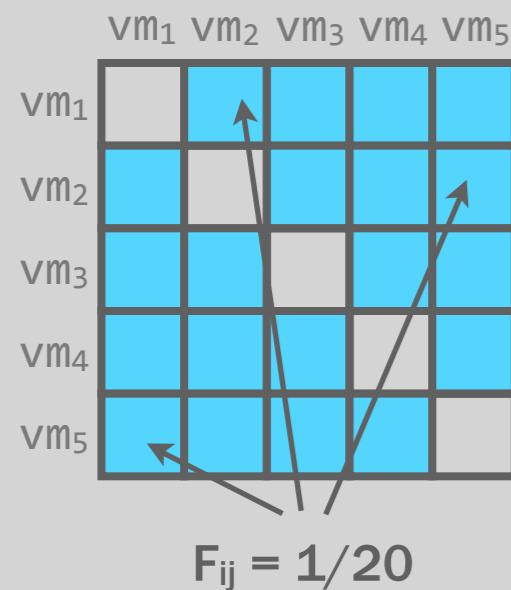


1. let F_{ij} = fraction of tenant traffic sent from VM_i to VM_j

SPATIAL VARIABILITY

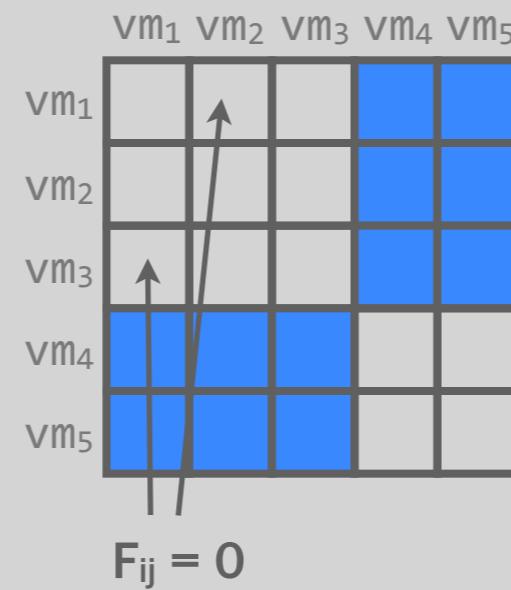
how do we quantify spatial variability?

no spatial variability

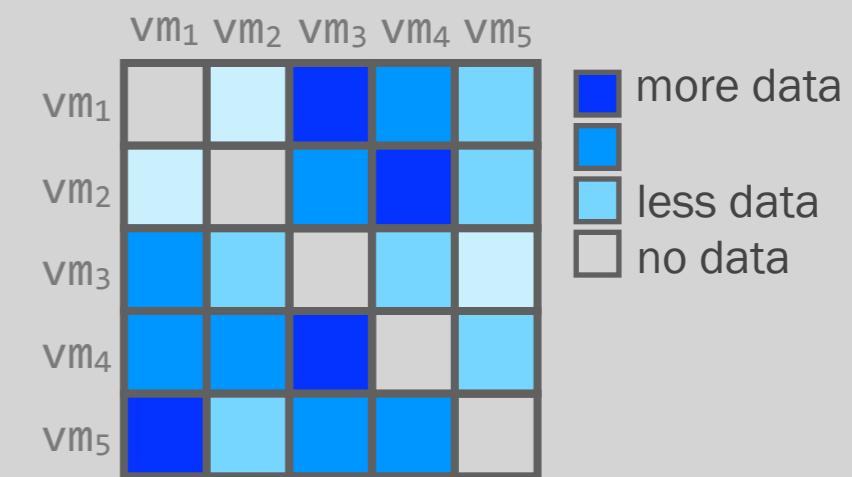


all F_{ij} values are equal

some spatial variability



high spatial variability

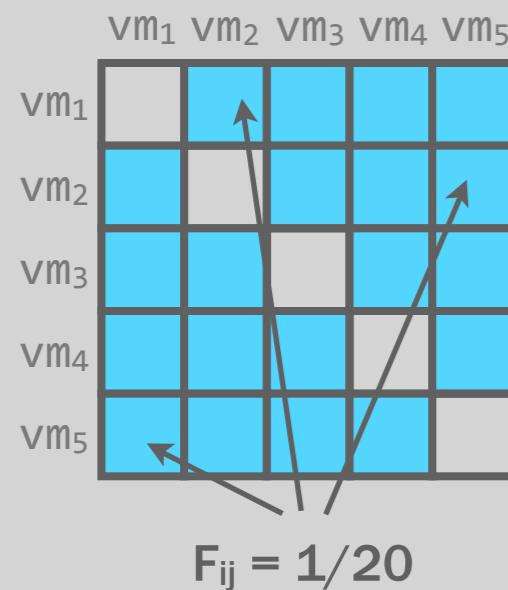


1. let F_{ij} = fraction of tenant traffic sent from VM_i to VM_j

SPATIAL VARIABILITY

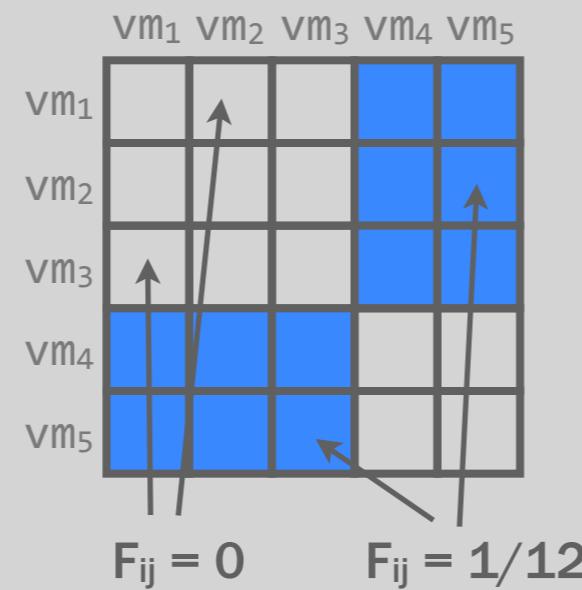
how do we quantify spatial variability?

no spatial variability

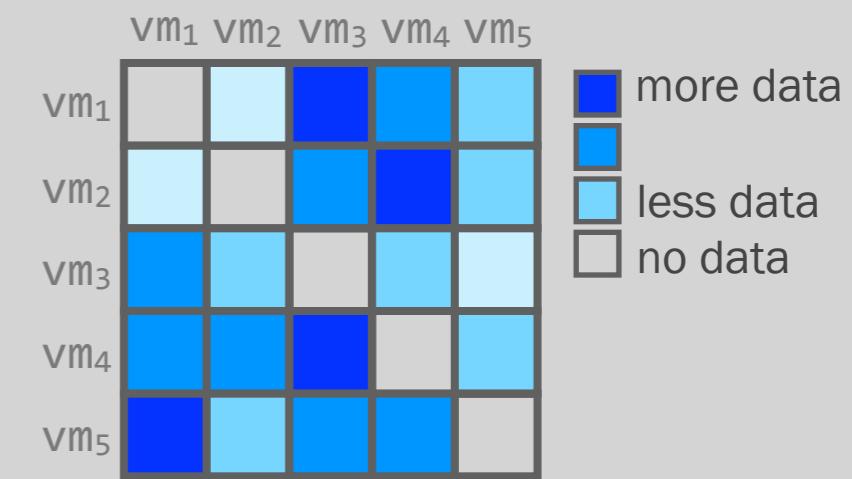


all F_{ij} values are equal

some spatial variability



high spatial variability

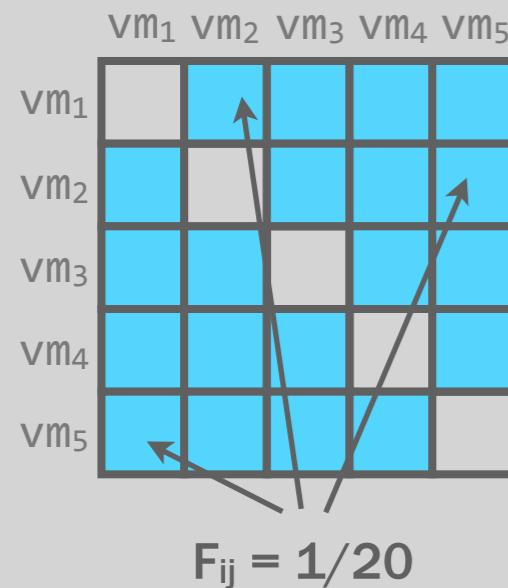


1. let F_{ij} = fraction of tenant traffic sent from VM_i to VM_j

SPATIAL VARIABILITY

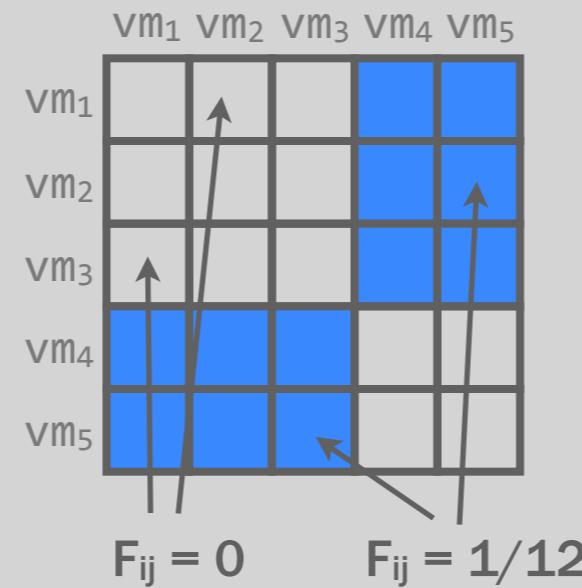
how do we quantify spatial variability?

no spatial variability



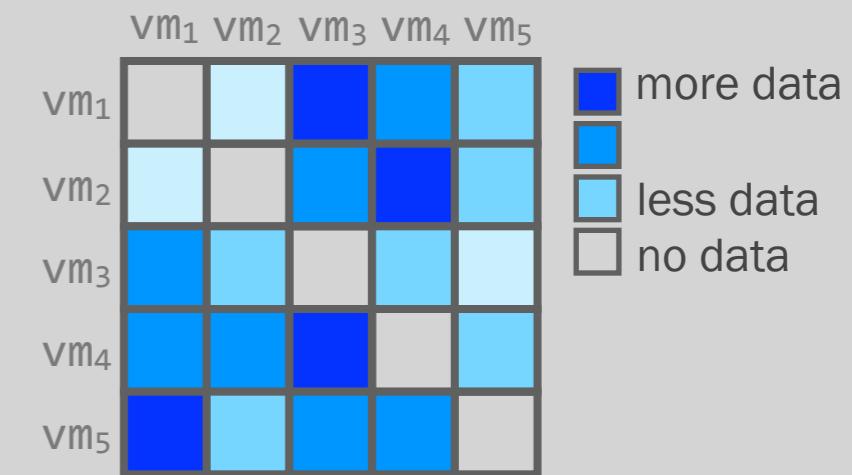
all F_{ij} values are equal

some spatial variability



two distinct F_{ij} values

high spatial variability

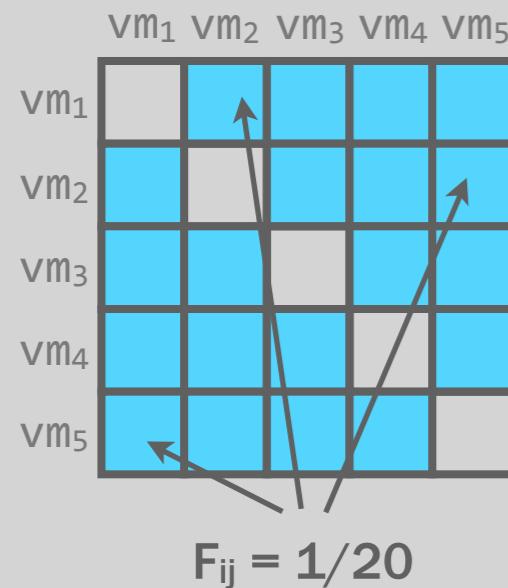


1. let F_{ij} = fraction of tenant traffic sent from VM_i to VM_j

SPATIAL VARIABILITY

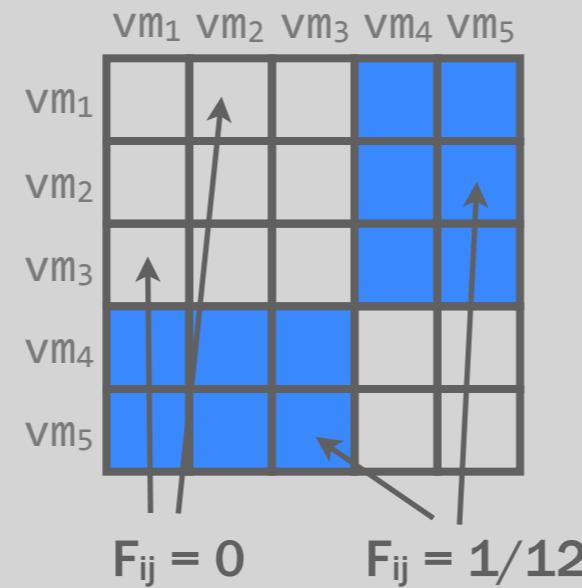
how do we quantify spatial variability?

no spatial variability

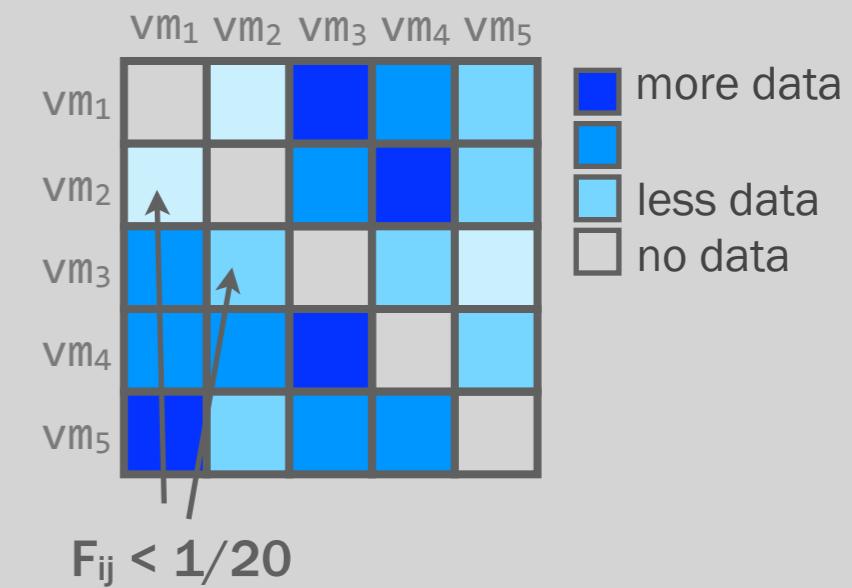


all F_{ij} values are equal

some spatial variability



high spatial variability

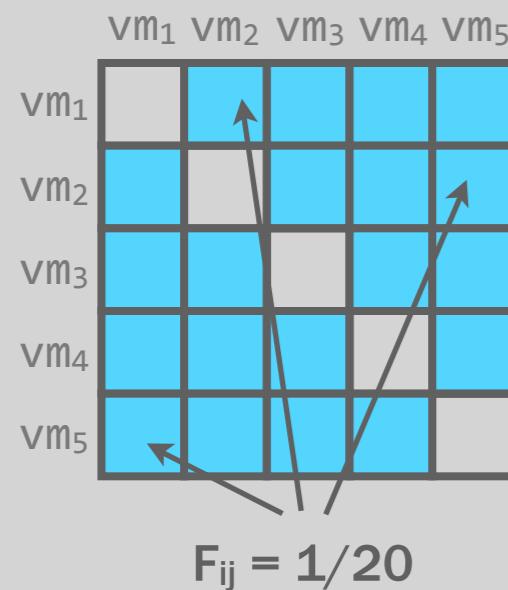


1. let F_{ij} = fraction of tenant traffic sent from VM_i to VM_j

SPATIAL VARIABILITY

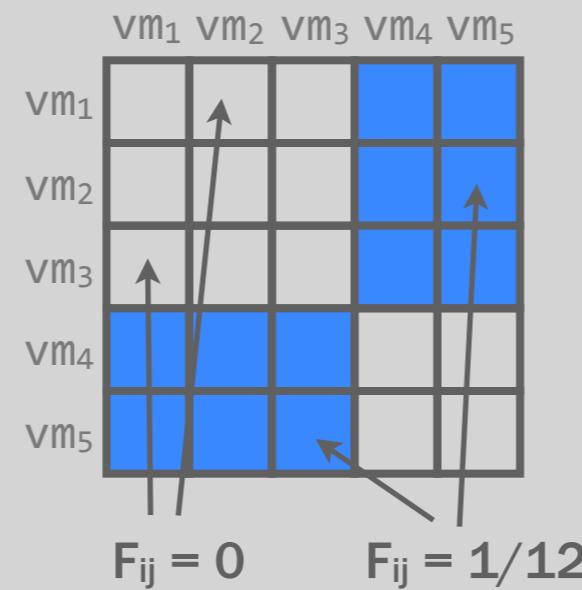
how do we quantify spatial variability?

no spatial variability



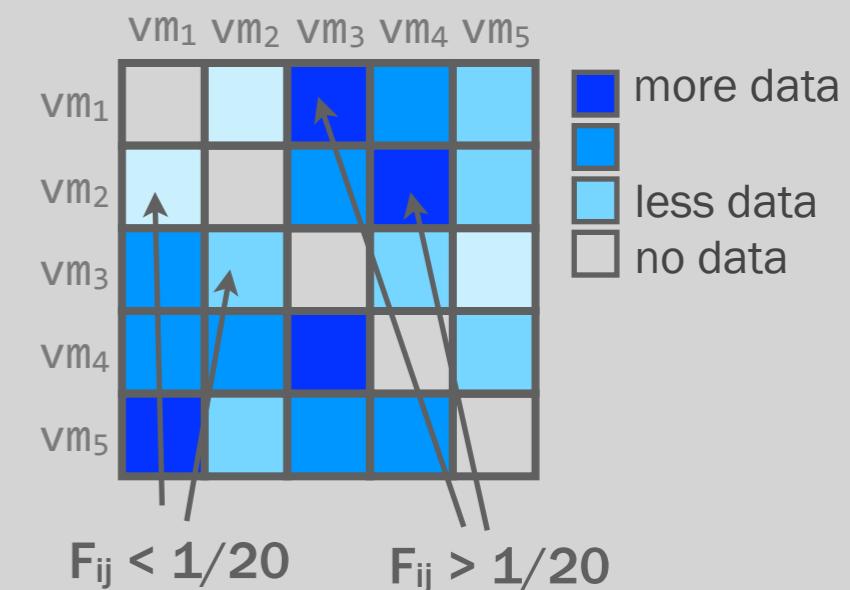
all F_{ij} values are equal

some spatial variability



two distinct F_{ij} values

high spatial variability

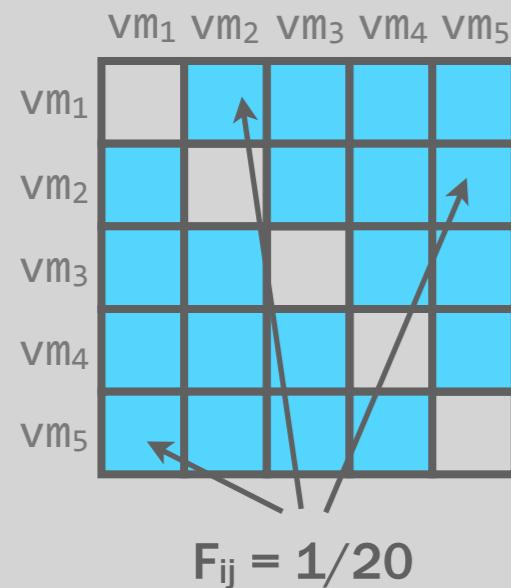


1. let F_{ij} = fraction of tenant traffic sent from VM_i to VM_j

SPATIAL VARIABILITY

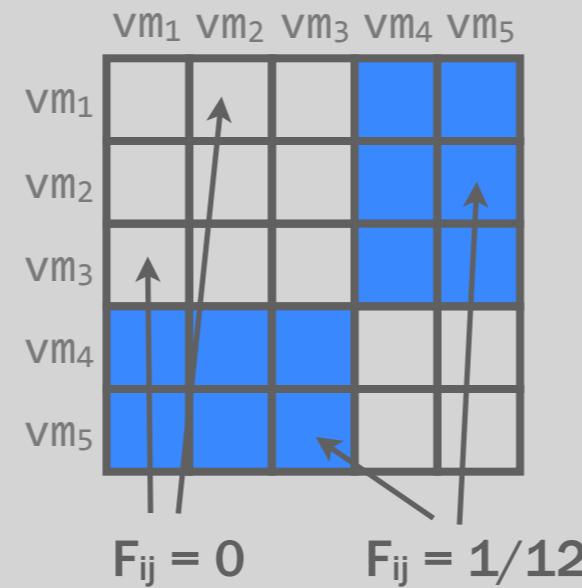
how do we quantify spatial variability?

no spatial variability

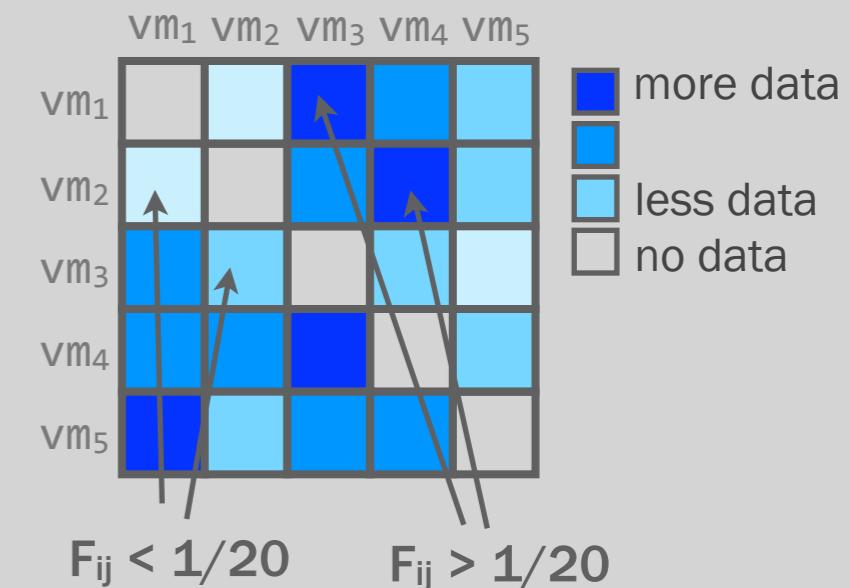


all F_{ij} values are equal

some spatial variability



high spatial variability

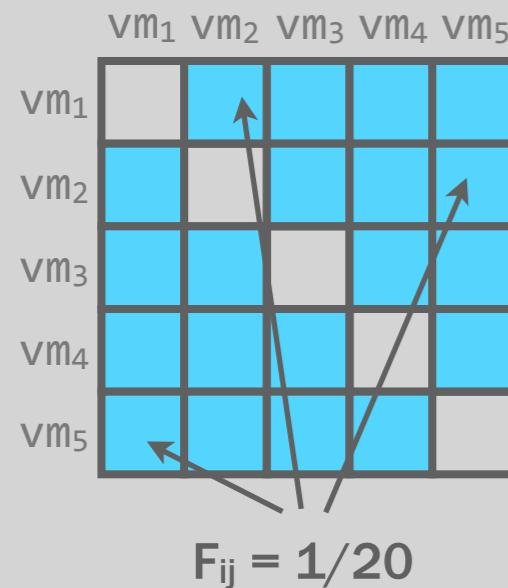


1. let F_{ij} = fraction of tenant traffic sent from VM_i to VM_j

SPATIAL VARIABILITY

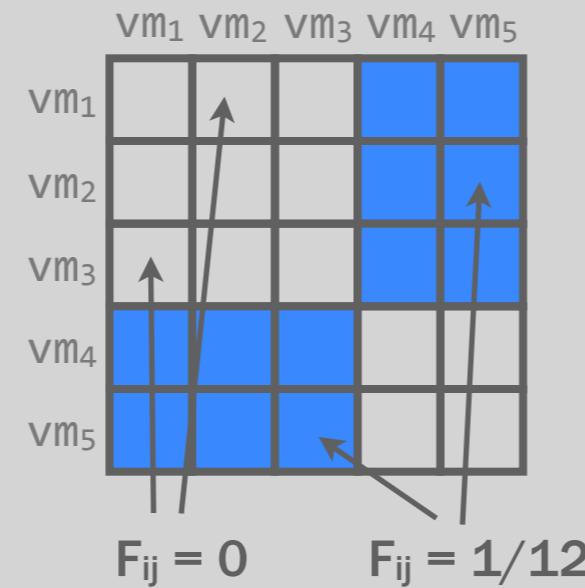
how do we quantify spatial variability?

no spatial variability



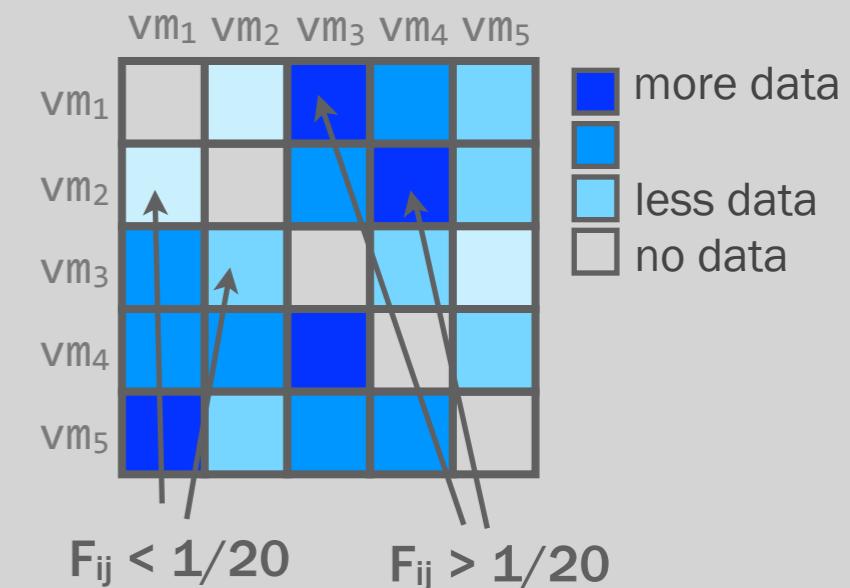
all F_{ij} values are equal

some spatial variability



two distinct F_{ij} values

high spatial variability



many distinct F_{ij} values

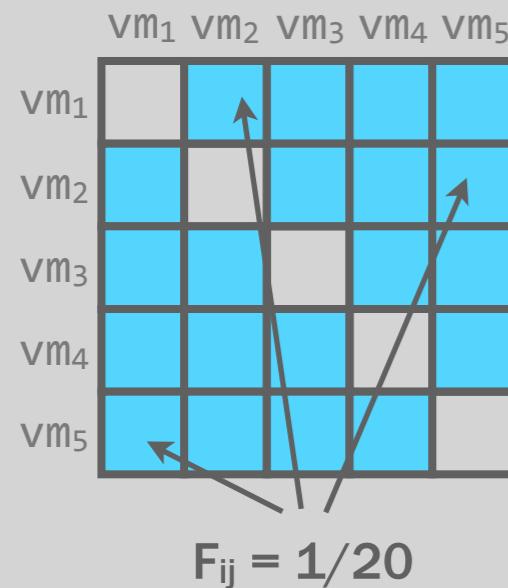
1. let F_{ij} = fraction of tenant traffic sent from VM_i to VM_j

2. calculate the coefficient of variation (σ/μ) of the F_{ij} values

SPATIAL VARIABILITY

how do we quantify spatial variability?

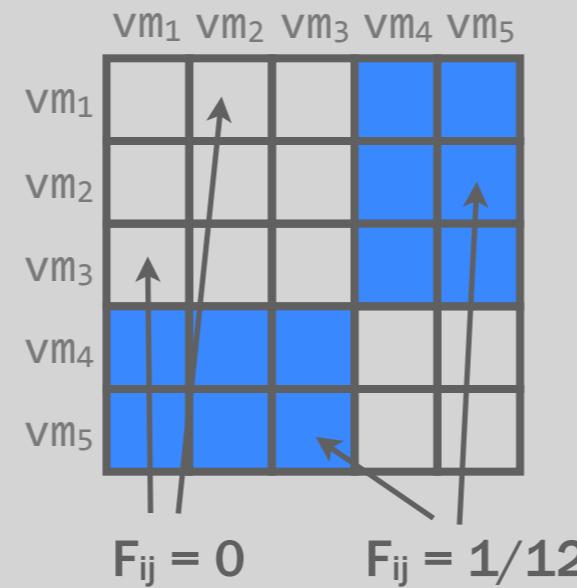
no spatial variability



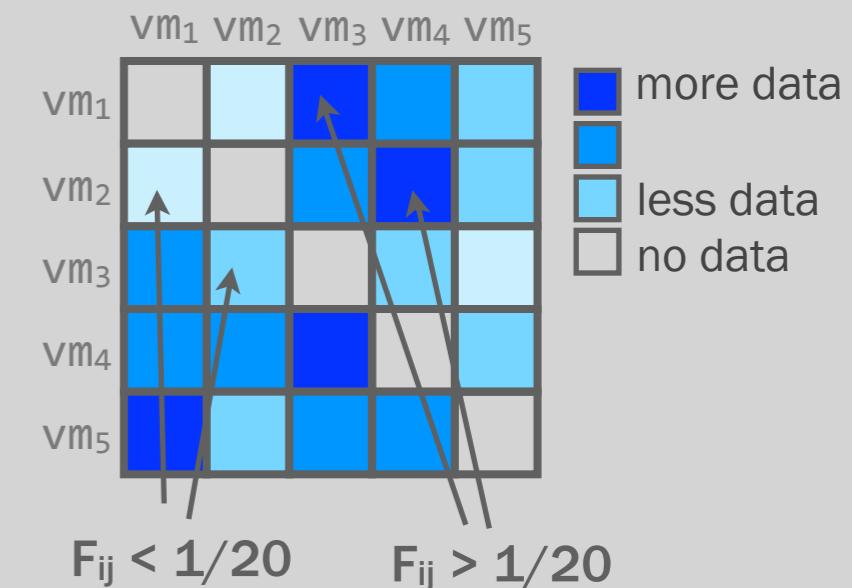
all F_{ij} values are equal

$$\Rightarrow \sigma(F_{ij})/\mu(F_{ij}) = 0$$

some spatial variability



high spatial variability



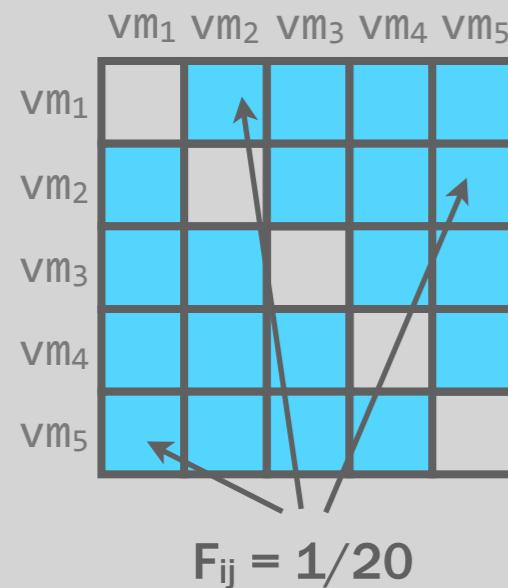
1. let F_{ij} = fraction of tenant traffic sent from VM_i to VM_j

2. calculate the coefficient of variation (σ/μ) of the F_{ij} values

SPATIAL VARIABILITY

how do we quantify spatial variability?

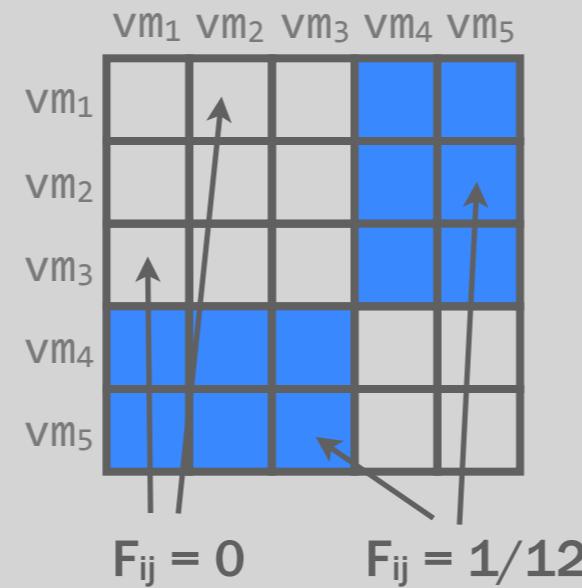
no spatial variability



all F_{ij} values are equal

$$\Rightarrow \sigma(F_{ij})/\mu(F_{ij}) = 0$$

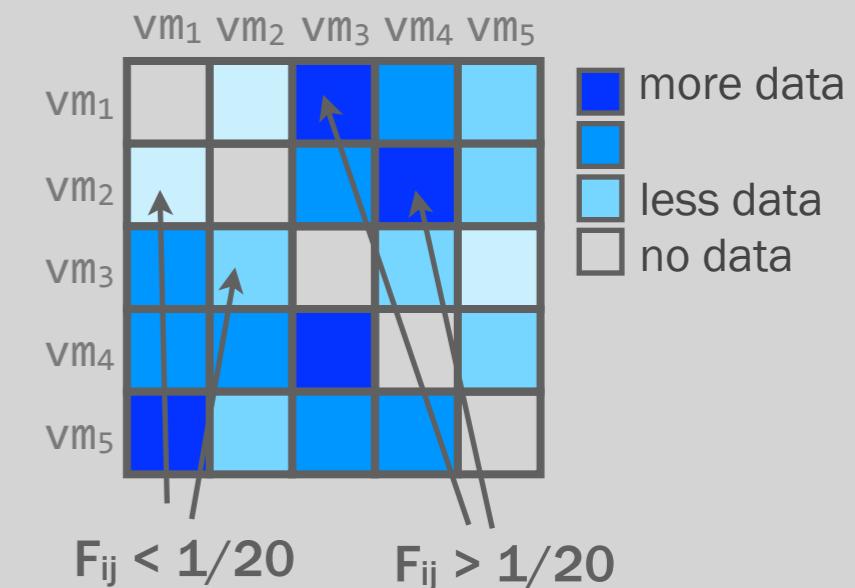
some spatial variability



two distinct F_{ij} values

$$\Rightarrow \sigma(F_{ij})/\mu(F_{ij}) = 1^*$$

high spatial variability



many distinct F_{ij} values

- more data
- less data
- no data

1. let F_{ij} = fraction of tenant traffic sent from VM_i to VM_j

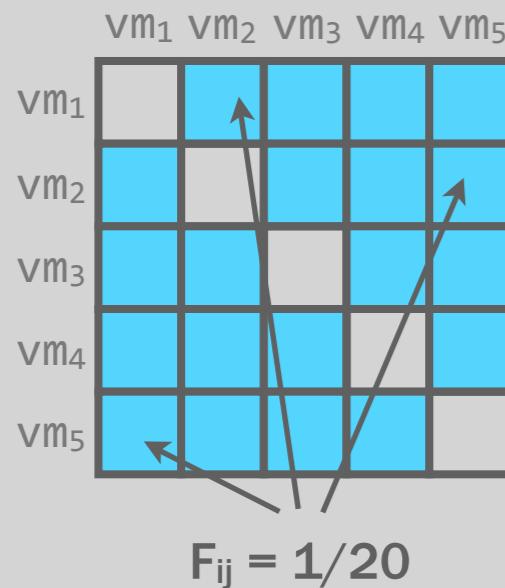
2. calculate the coefficient of variation (σ/μ) of the F_{ij} values

*see paper for this result

SPATIAL VARIABILITY

how do we quantify spatial variability?

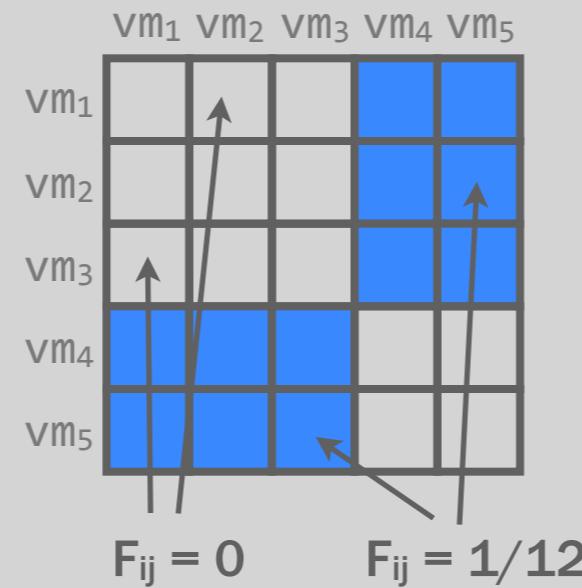
no spatial variability



all F_{ij} values are equal

$$\Rightarrow \sigma(F_{ij})/\mu(F_{ij}) = 0$$

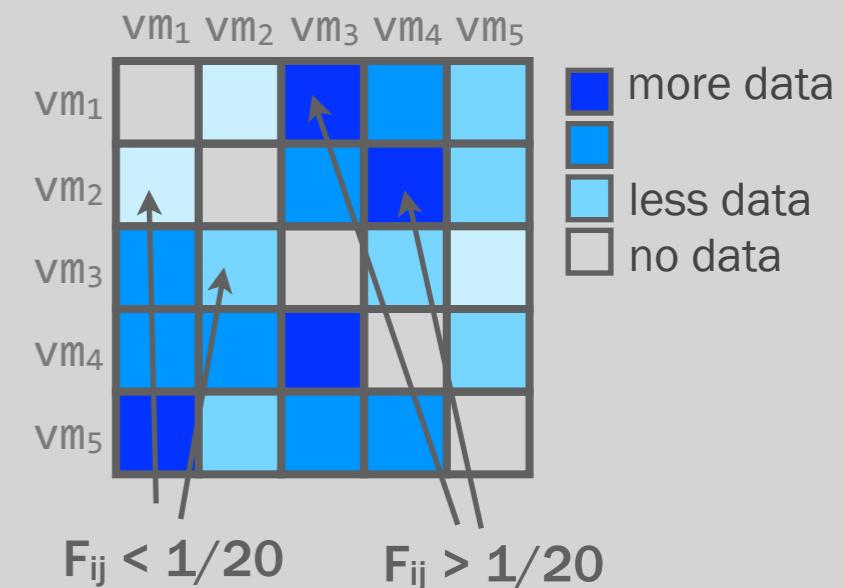
some spatial variability



two distinct F_{ij} values

$$\Rightarrow \sigma(F_{ij})/\mu(F_{ij}) = 1^*$$

high spatial variability



many distinct F_{ij} values

$$\Rightarrow \sigma(F_{ij})/\mu(F_{ij}) > 1$$

1. let F_{ij} = fraction of tenant traffic sent from VM_i to VM_j

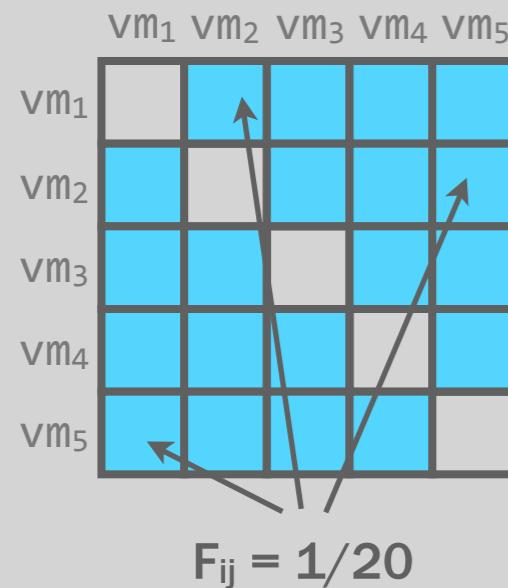
2. calculate the coefficient of variation (σ/μ) of the F_{ij} values

*see paper for this result

SPATIAL VARIABILITY

how do we quantify spatial variability?

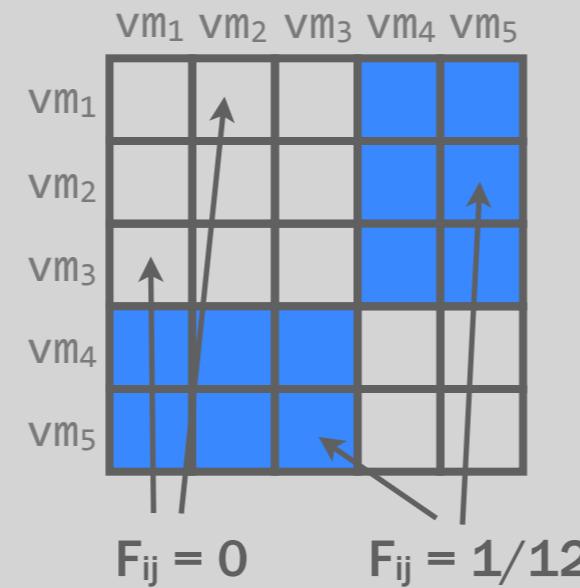
no spatial variability



all F_{ij} values are equal

$$\Rightarrow \sigma(F_{ij})/\mu(F_{ij}) = 0$$

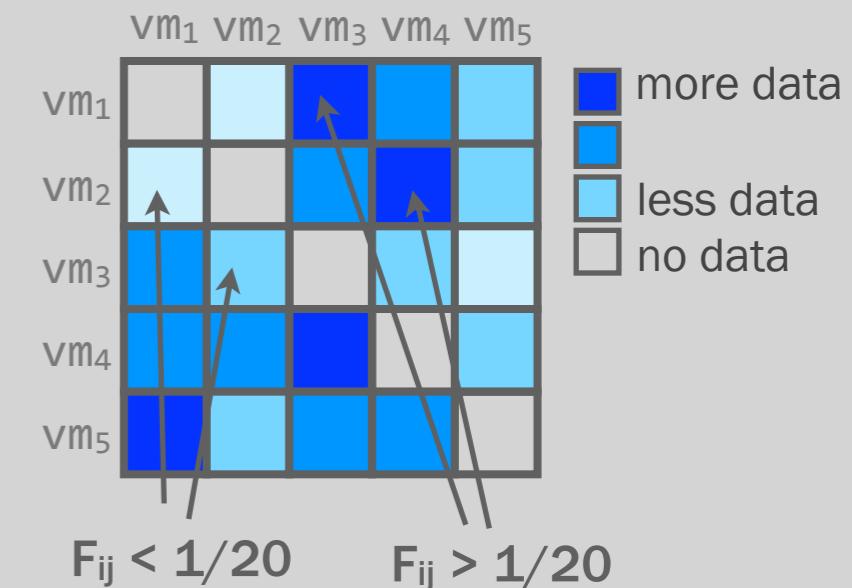
some spatial variability



two distinct F_{ij} values

$$\Rightarrow \sigma(F_{ij})/\mu(F_{ij}) = 1^*$$

high spatial variability



many distinct F_{ij} values

$$\Rightarrow \sigma(F_{ij})/\mu(F_{ij}) > 1$$

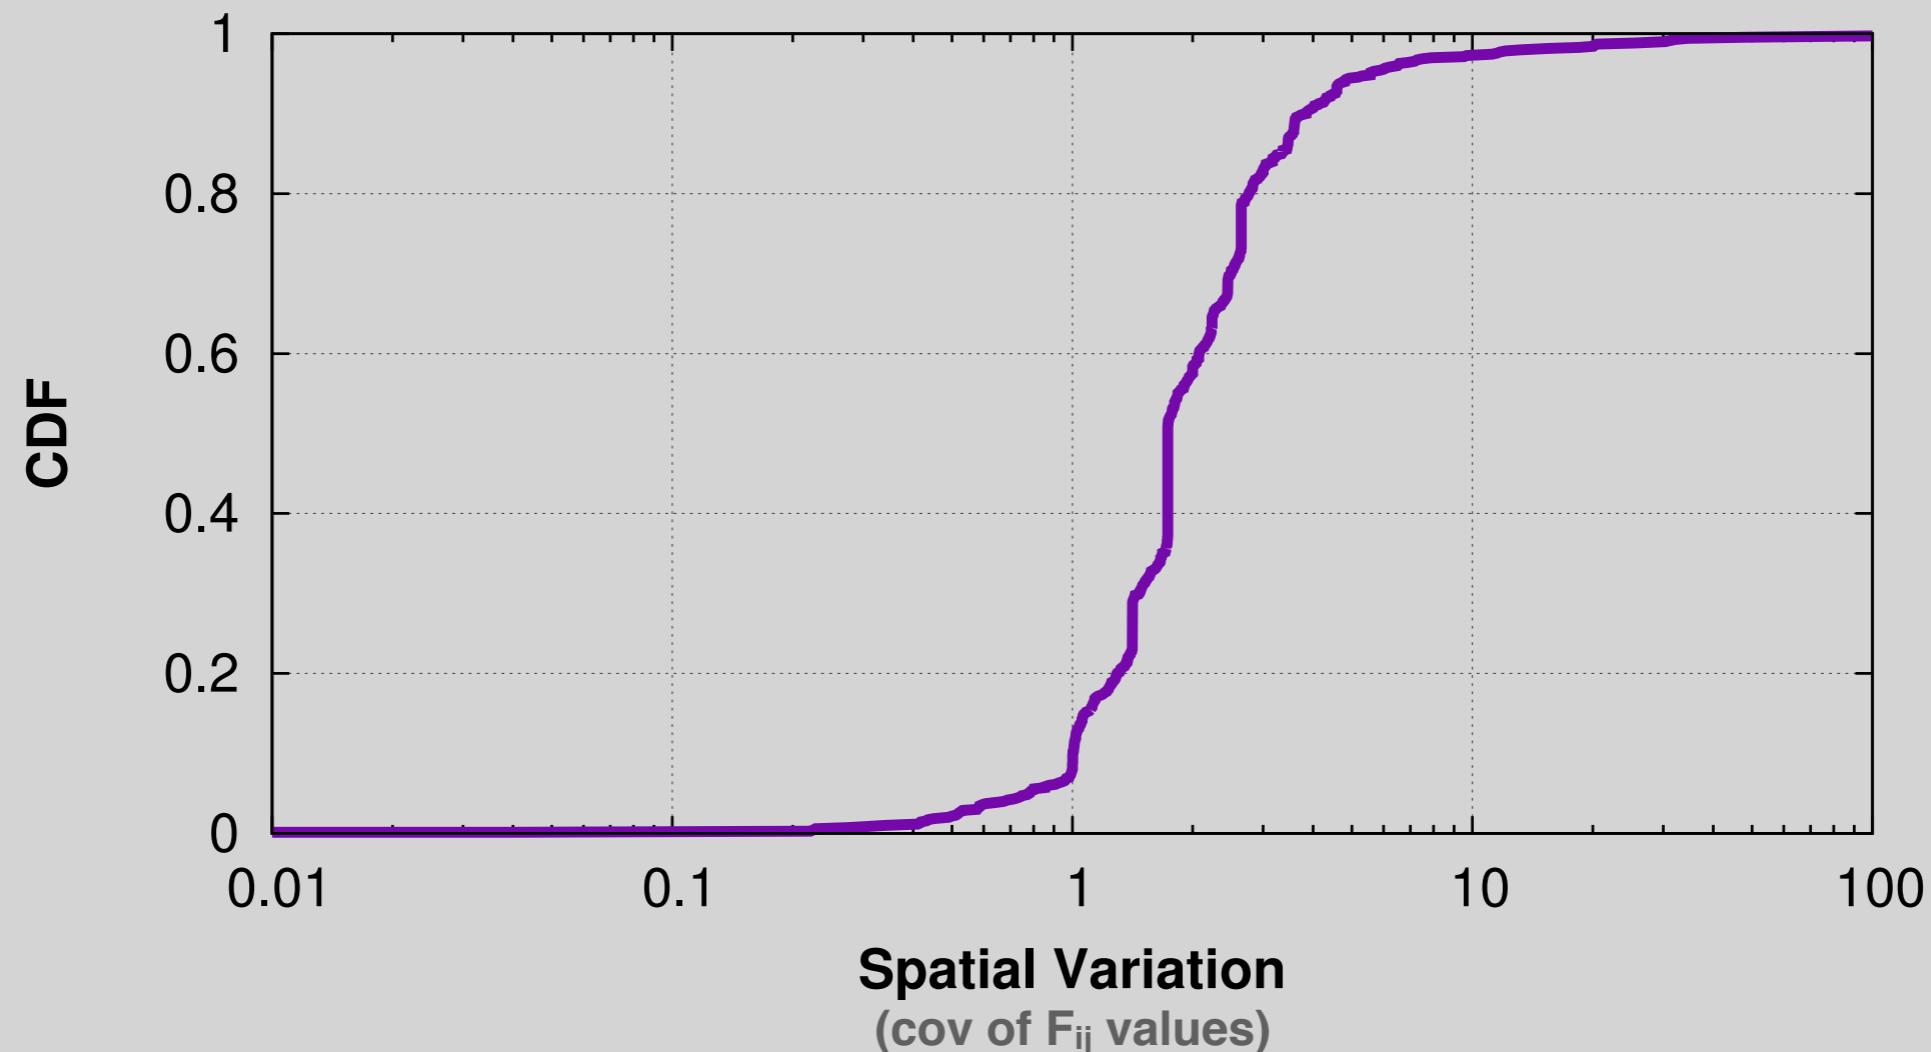
1. let F_{ij} = fraction of tenant traffic sent from VM_i to VM_j

2. calculate the coefficient of variation (σ/μ) of the F_{ij} values

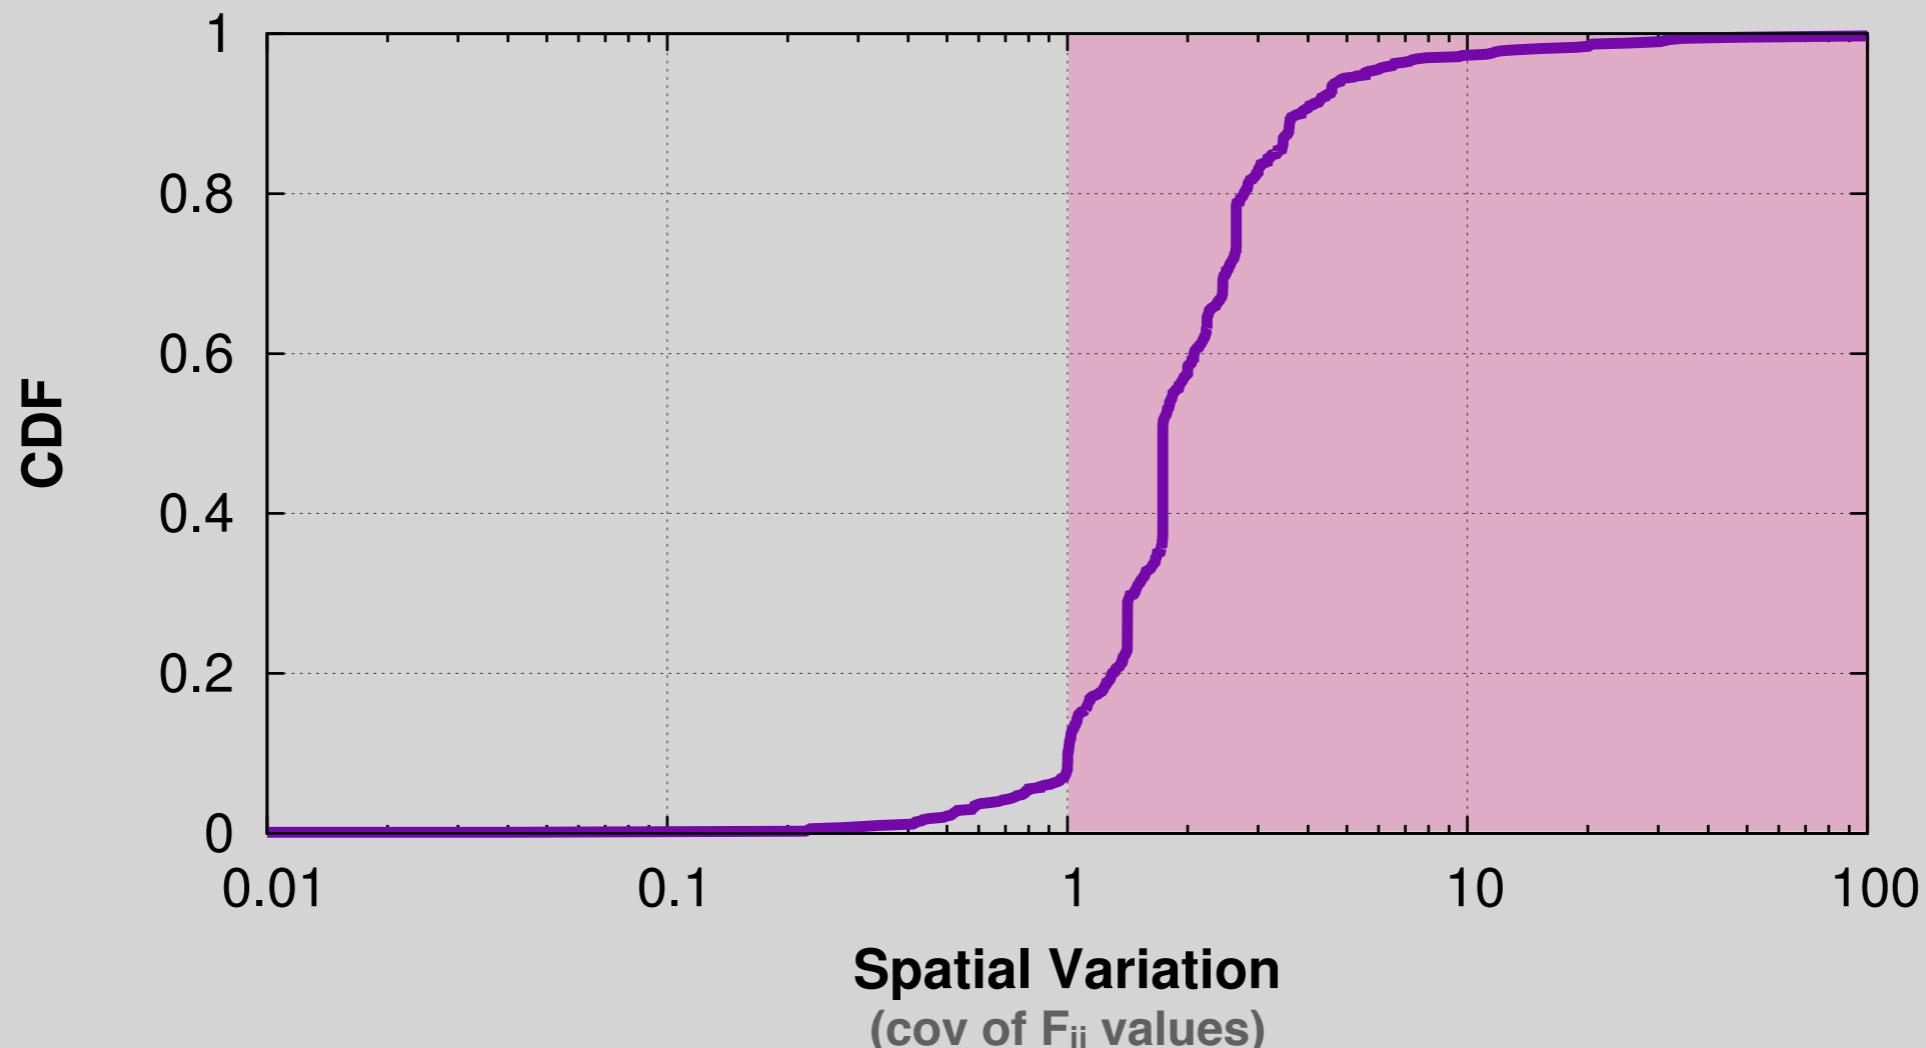
the higher the cov, the higher the spatial variation

*see paper for this result

SPATIAL VARIABILITY

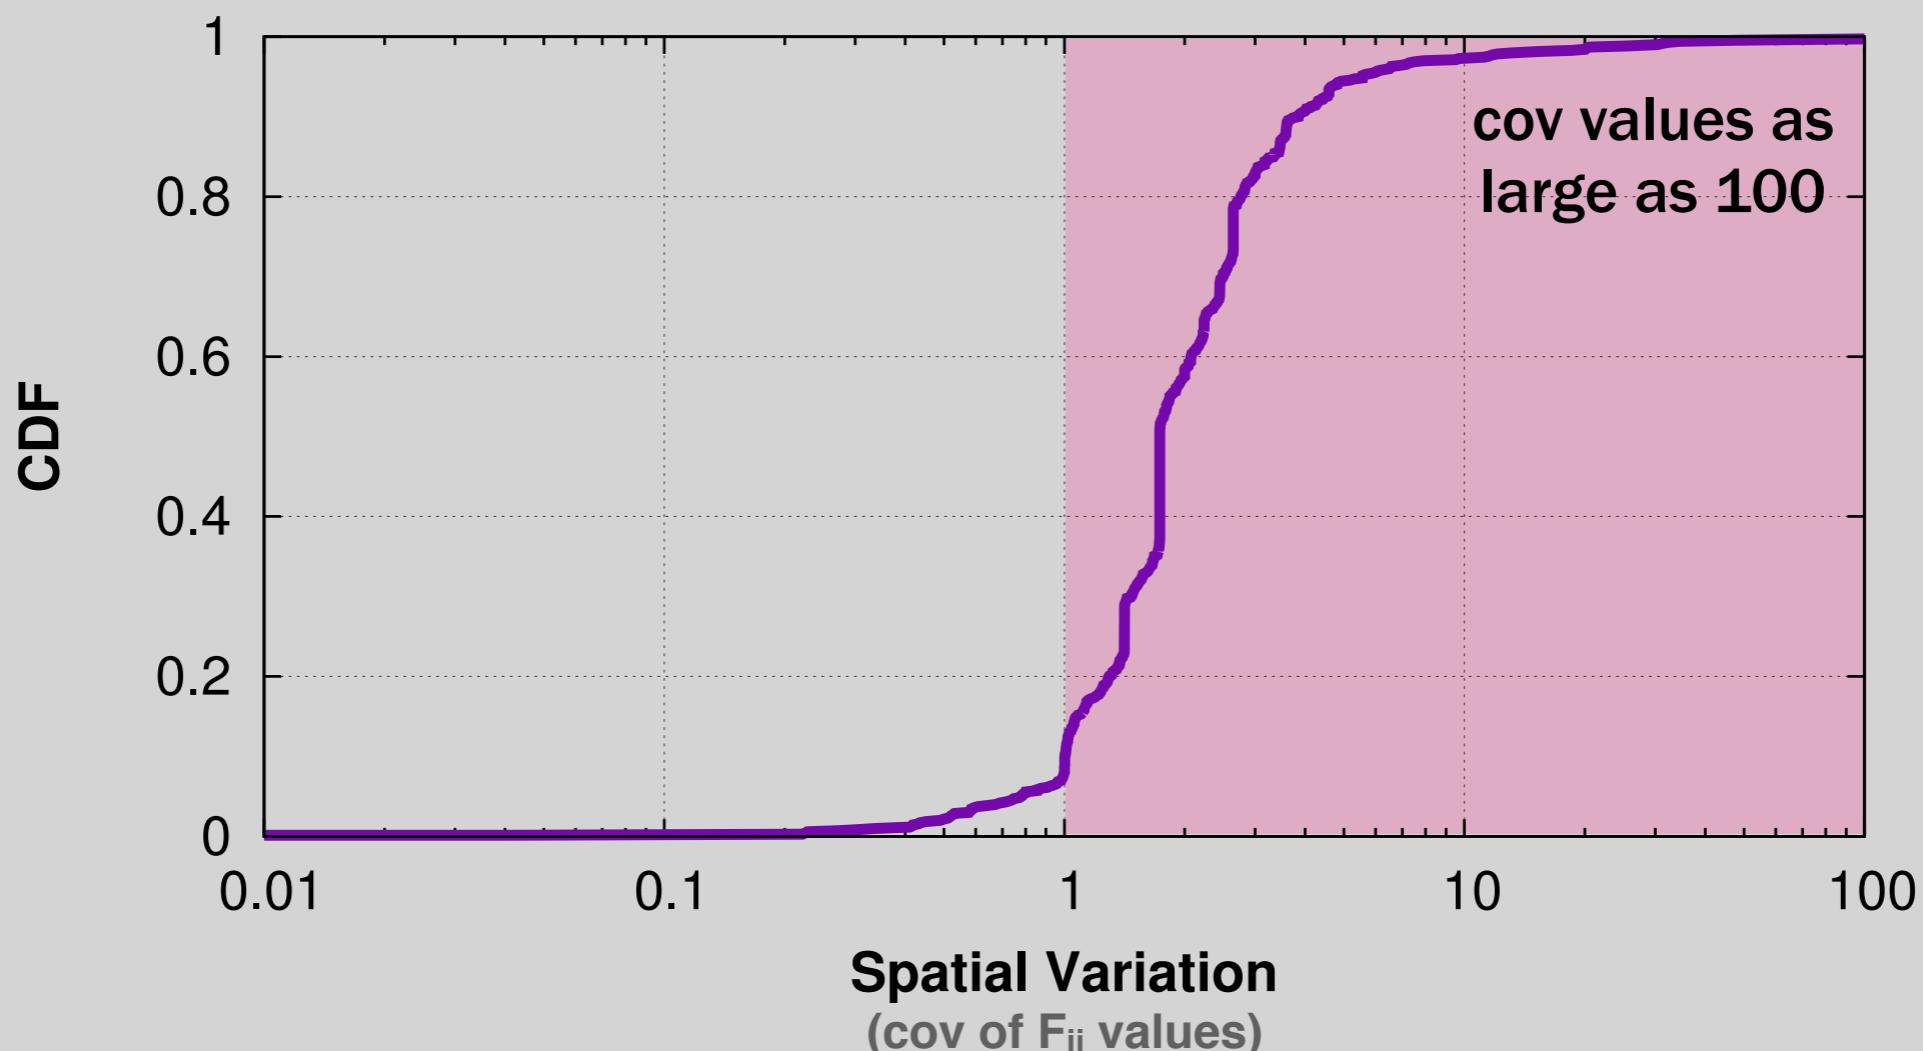


SPATIAL VARIABILITY



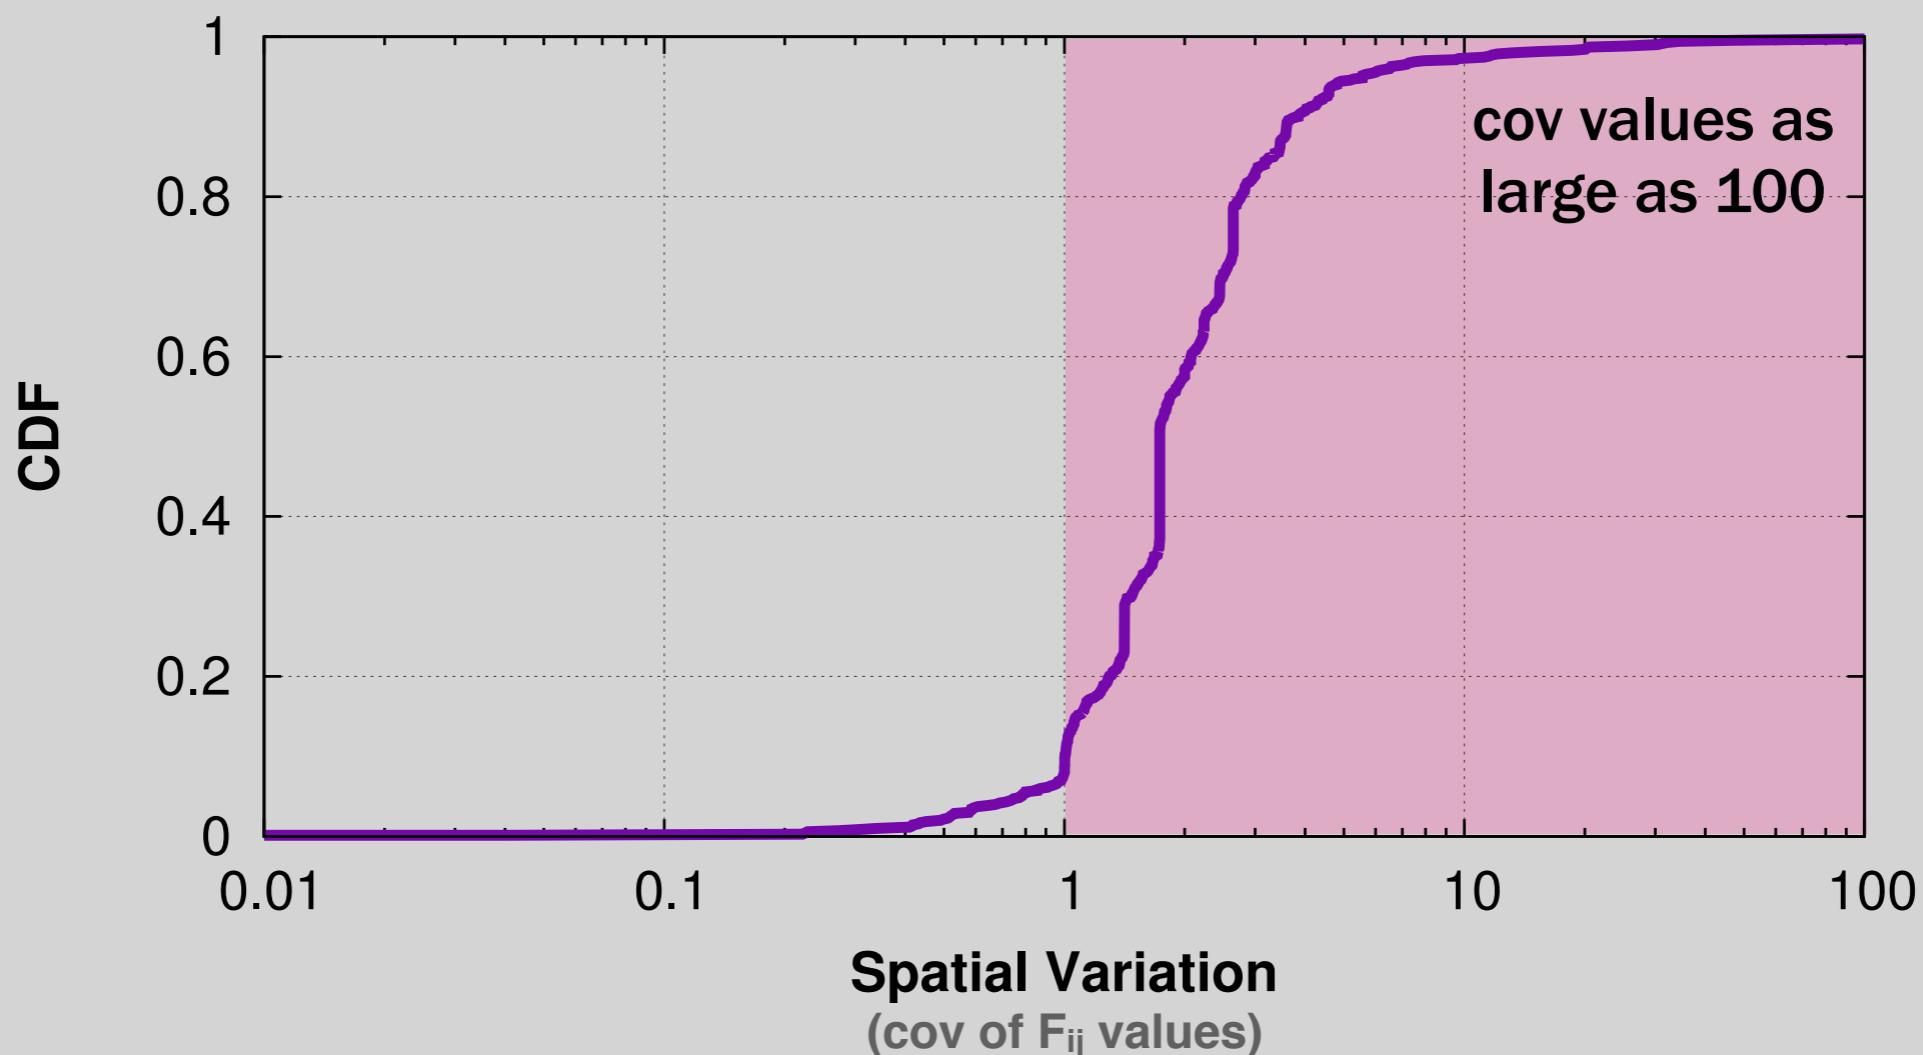
applications exhibit **high** spatial variability

SPATIAL VARIABILITY



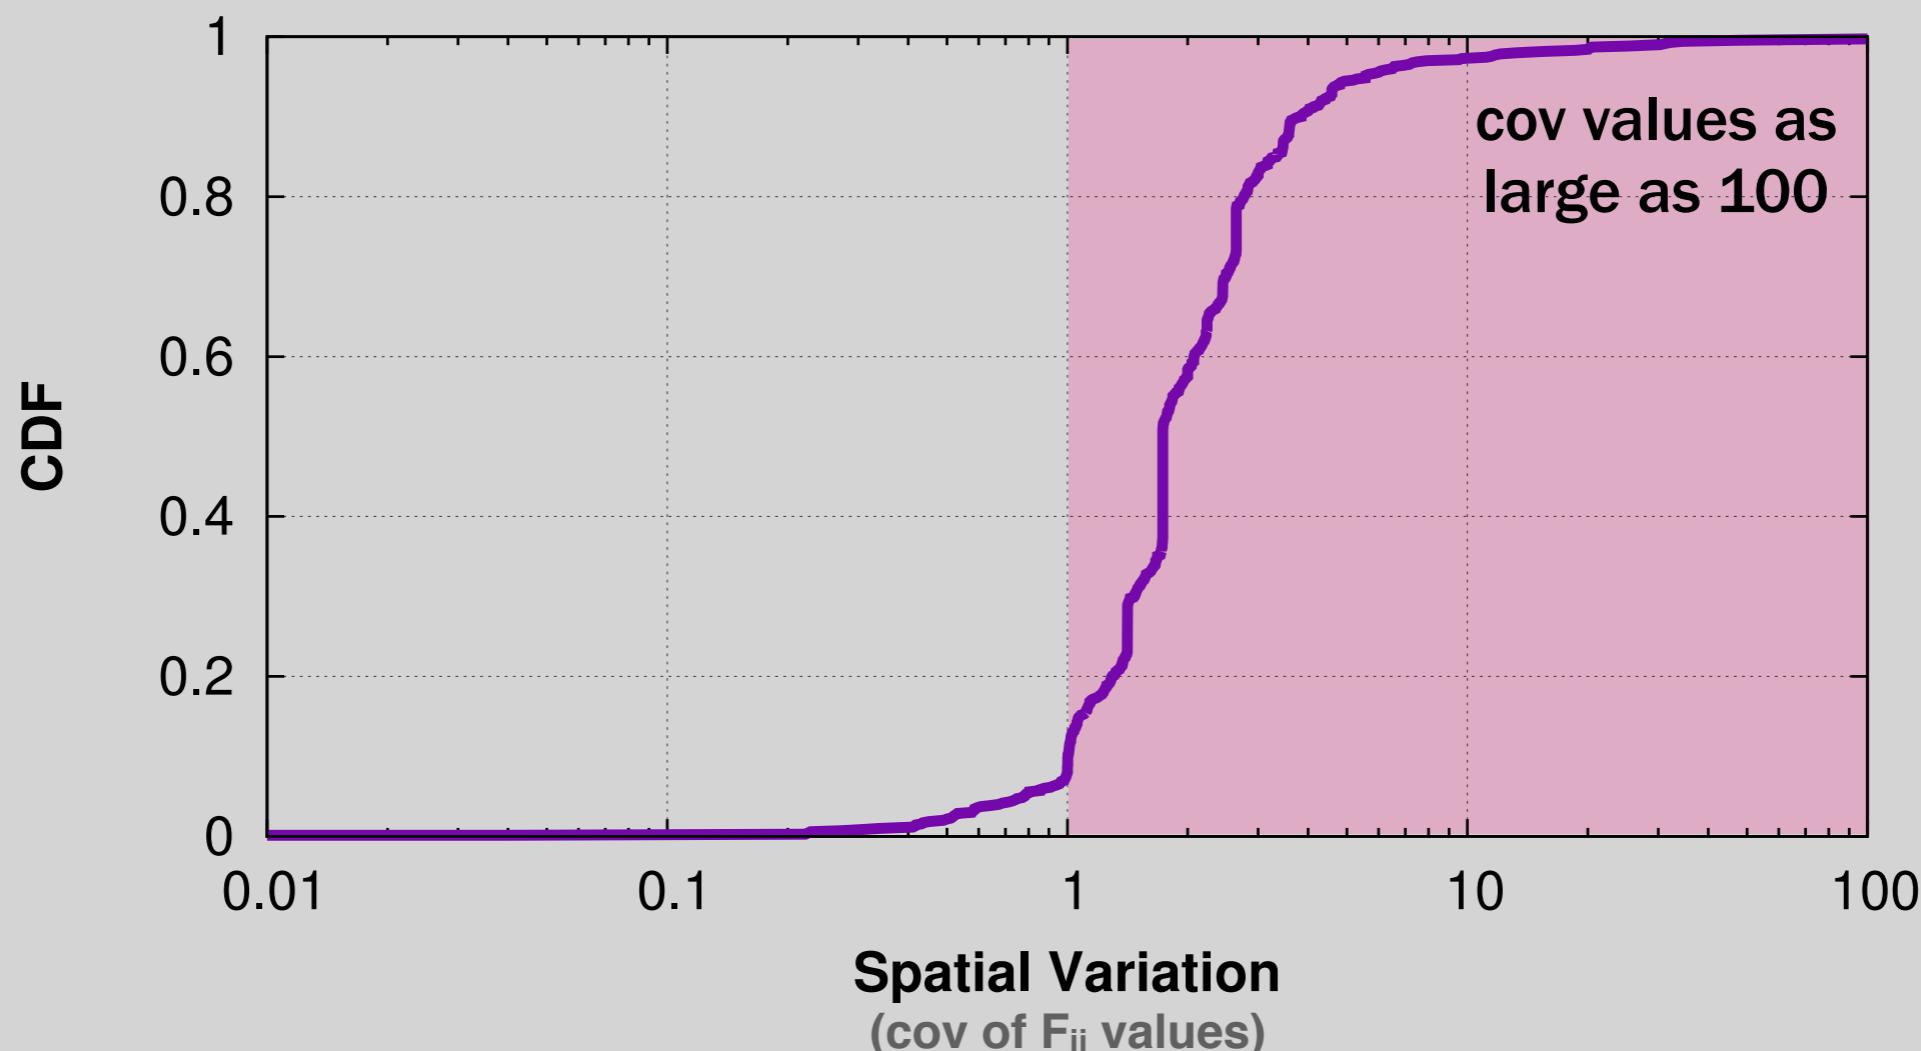
applications exhibit **high** spatial variability

SPATIAL VARIABILITY



applications exhibit **high** spatial variability
(the same result holds for temporal variability, see paper)

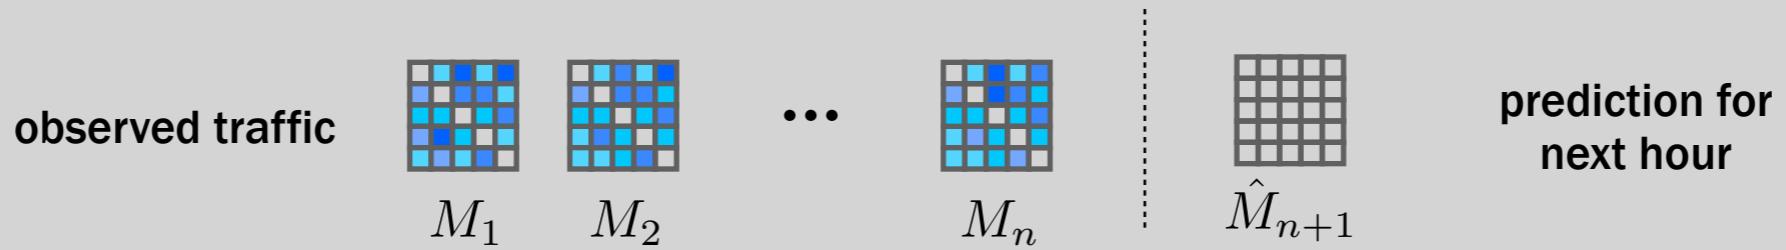
SPATIAL VARIABILITY



applications exhibit **high** spatial variability
(the same result holds for temporal variability, see paper)

customers cannot model such variability themselves,
and existing models do not capture it

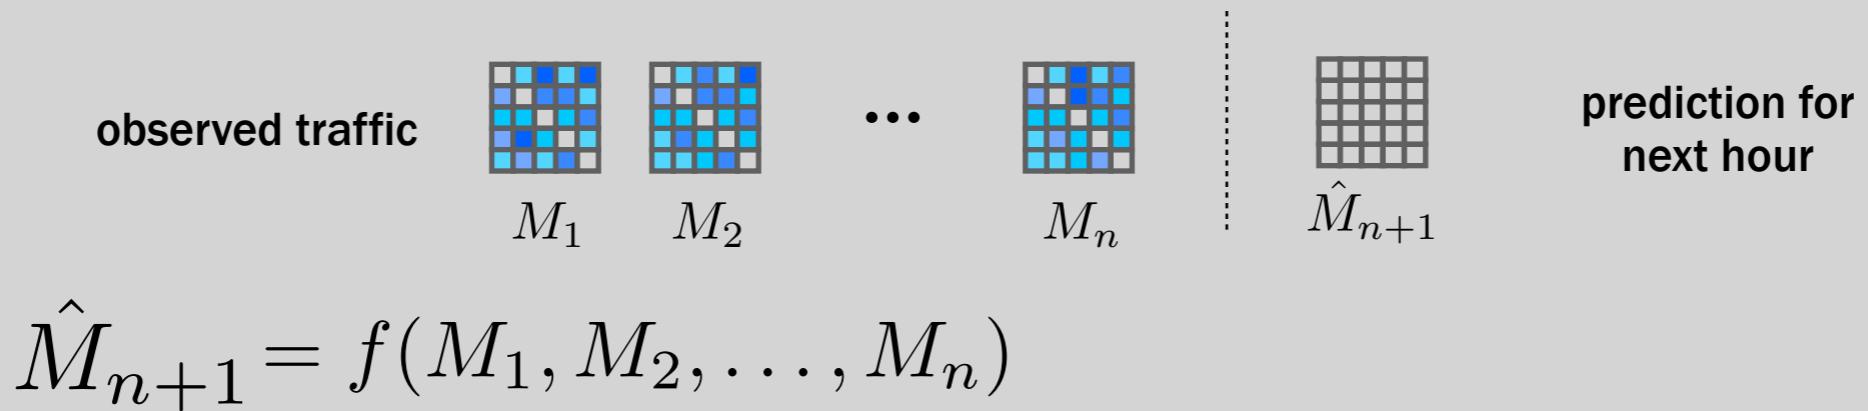
CICADA'S ALGORITHM



Cicada's prediction algorithm draws inspiration from the following works:

- [1] Wang, et al. "COPE: Traffic Engineering in Dynamic Networks". SIGCOMM 2006.
- [2] Herbster and Warmuth. "Tracking the Best Expert". Machine Learning Vol 32, 1998.

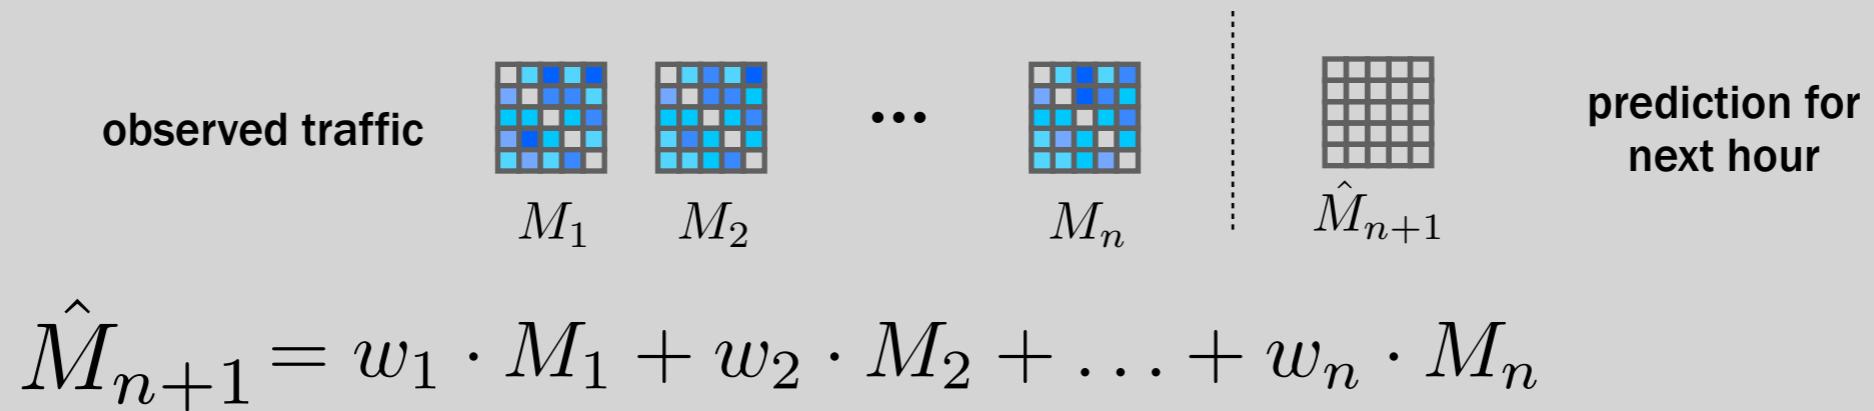
CICADA'S ALGORITHM



Cicada's prediction algorithm draws inspiration from the following works:

- [1] Wang, et al. "COPE: Traffic Engineering in Dynamic Networks". SIGCOMM 2006.
- [2] Herbster and Warmuth. "Tracking the Best Expert". Machine Learning Vol 32, 1998.

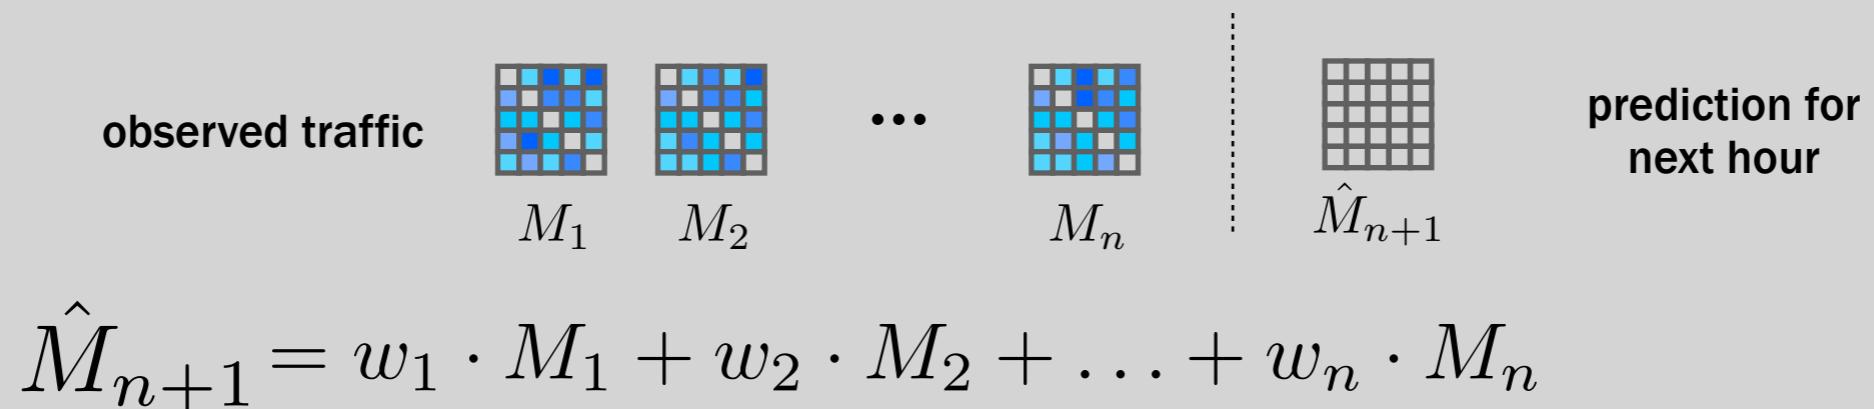
CICADA'S ALGORITHM



Cicada's prediction algorithm draws inspiration from the following works:

- [1] Wang, et al. "COPE: Traffic Engineering in Dynamic Networks". SIGCOMM 2006.
- [2] Herbster and Warmuth. "Tracking the Best Expert". Machine Learning Vol 32, 1998.

CICADA'S ALGORITHM



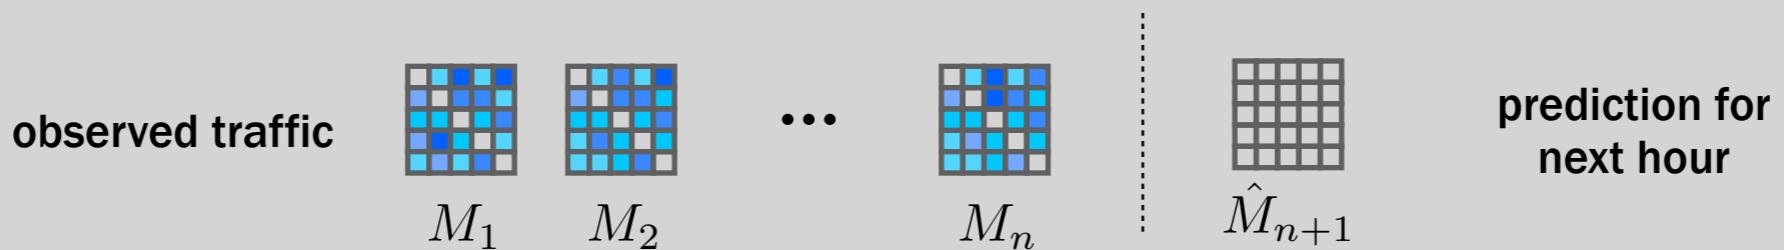
instead of setting the weights beforehand, cicada learns the weights **online**; they are updated with every prediction based on errors made in previous predictions

Cicada's prediction algorithm draws inspiration from the following works:

[1] Wang, et al. "COPE: Traffic Engineering in Dynamic Networks". SIGCOMM 2006.

[2] Herbster and Warmuth. "Tracking the Best Expert". Machine Learning Vol 32, 1998.

CICADA'S ALGORITHM



$$\hat{M}_{n+1} = w_1(n) \cdot M_1 + w_2(n) \cdot M_2 + \dots + w_n(n) \cdot M_n$$

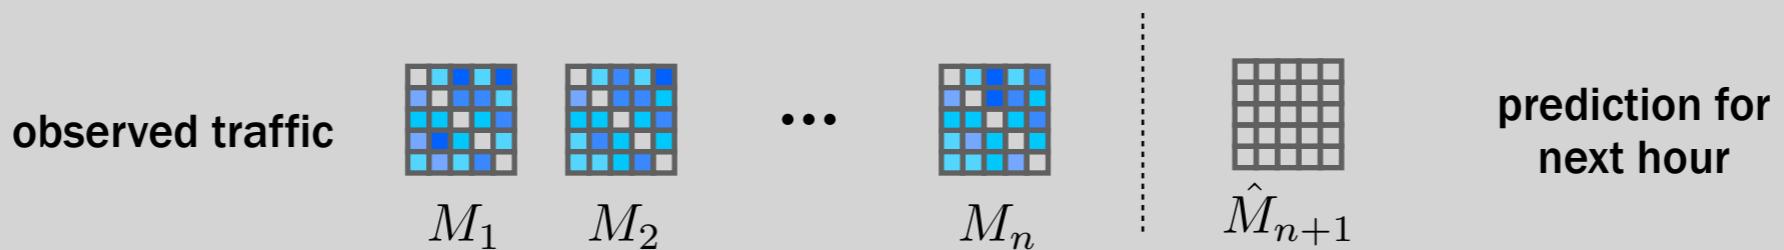
instead of setting the weights beforehand, cicada learns the weights **online**; they are updated with every prediction based on errors made in previous predictions

Cicada's prediction algorithm draws inspiration from the following works:

[1] Wang, et al. "COPE: Traffic Engineering in Dynamic Networks". SIGCOMM 2006.

[2] Herbster and Warmuth. "Tracking the Best Expert". Machine Learning Vol 32, 1998.

CICADA'S ALGORITHM



$$\hat{M}_{n+1} = w_1(n) \cdot M_1 + w_2(n) \cdot M_2 + \dots + w_n(n) \cdot M_n$$

instead of setting the weights beforehand, cicada learns the weights **online**; they are updated with every prediction based on errors made in previous predictions

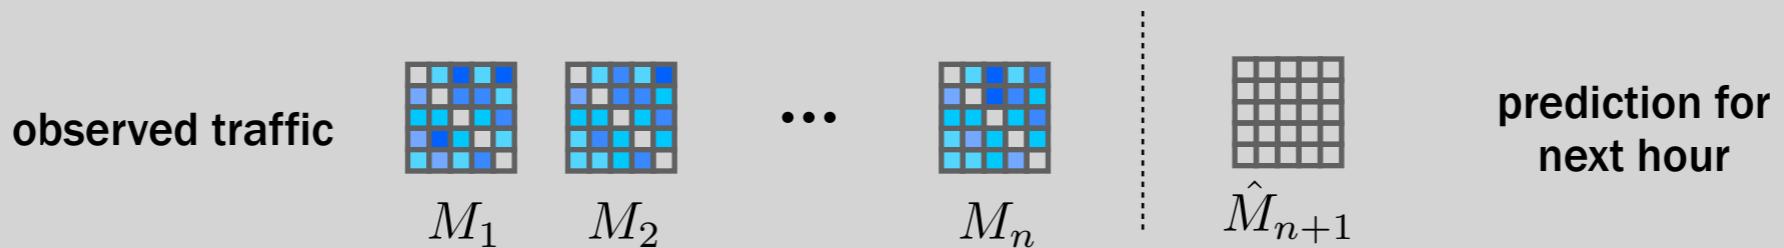
$$w_i(n+1) = w_i(n)$$

Cicada's prediction algorithm draws inspiration from the following works:

[1] Wang, et al. "COPE: Traffic Engineering in Dynamic Networks". SIGCOMM 2006.

[2] Herbster and Warmuth. "Tracking the Best Expert". Machine Learning Vol 32, 1998.

CICADA'S ALGORITHM



$$\hat{M}_{n+1} = w_1(n) \cdot M_1 + w_2(n) \cdot M_2 + \dots + w_n(n) \cdot M_n$$

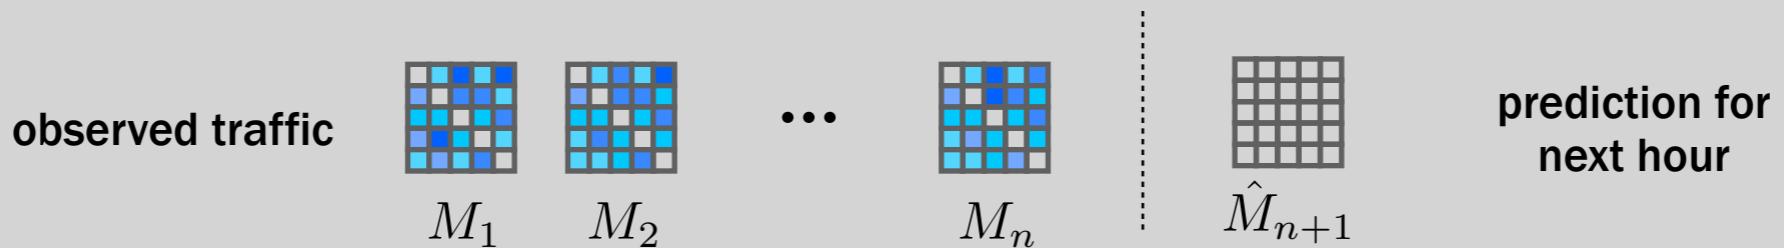
instead of setting the weights beforehand, cicada learns the weights **online**; they are updated with every prediction based on errors made in previous predictions

$$w_i(n+1) = w_i(n) \cdot e^{-L(i,n)}$$

loss function

Cicada's prediction algorithm draws inspiration from the following works:
[1] Wang, et al. "COPE: Traffic Engineering in Dynamic Networks". SIGCOMM 2006.
[2] Herbster and Warmuth. "Tracking the Best Expert". Machine Learning Vol 32, 1998.

CICADA'S ALGORITHM



$$\hat{M}_{n+1} = w_1(n) \cdot M_1 + w_2(n) \cdot M_2 + \dots + w_n(n) \cdot M_n$$

instead of setting the weights beforehand, cicada learns the weights **online**; they are updated with every prediction based on errors made in previous predictions

$$w_i(n+1) = w_i(n) \cdot e^{-L(i,n)}$$

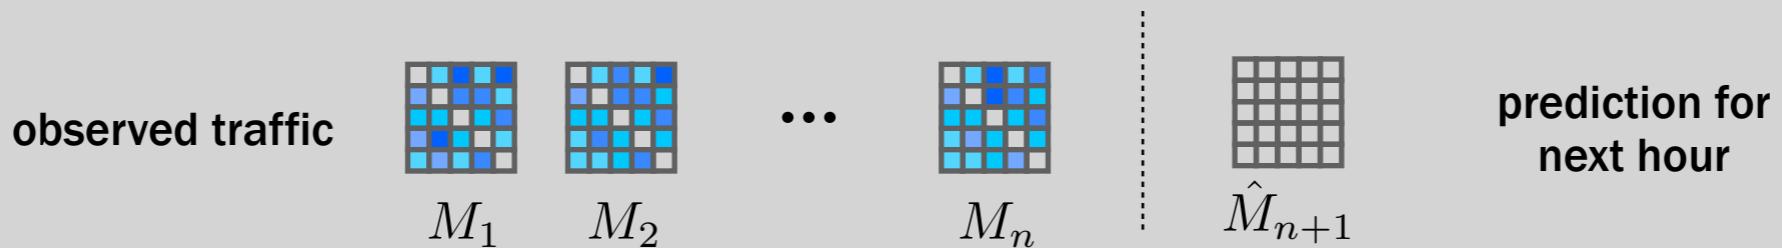
$$L(i, n) = \frac{\|\vec{M}_i - \vec{M}_n\|}{\|\vec{M}_n\|}$$

Cicada's prediction algorithm draws inspiration from the following works:

[1] Wang, et al. "COPE: Traffic Engineering in Dynamic Networks". SIGCOMM 2006.

[2] Herbster and Warmuth. "Tracking the Best Expert". Machine Learning Vol 32, 1998.

CICADA'S ALGORITHM



$$\hat{M}_{n+1} = w_1(n) \cdot M_1 + w_2(n) \cdot M_2 + \dots + w_n(n) \cdot M_n$$

instead of setting the weights beforehand, cicada learns the weights **online**; they are updated with every prediction based on errors made in previous predictions

$$w_i(n+1) = w_i(n) \cdot e^{-L(i,n)} \cdot \frac{1}{Z_{n+1}}$$

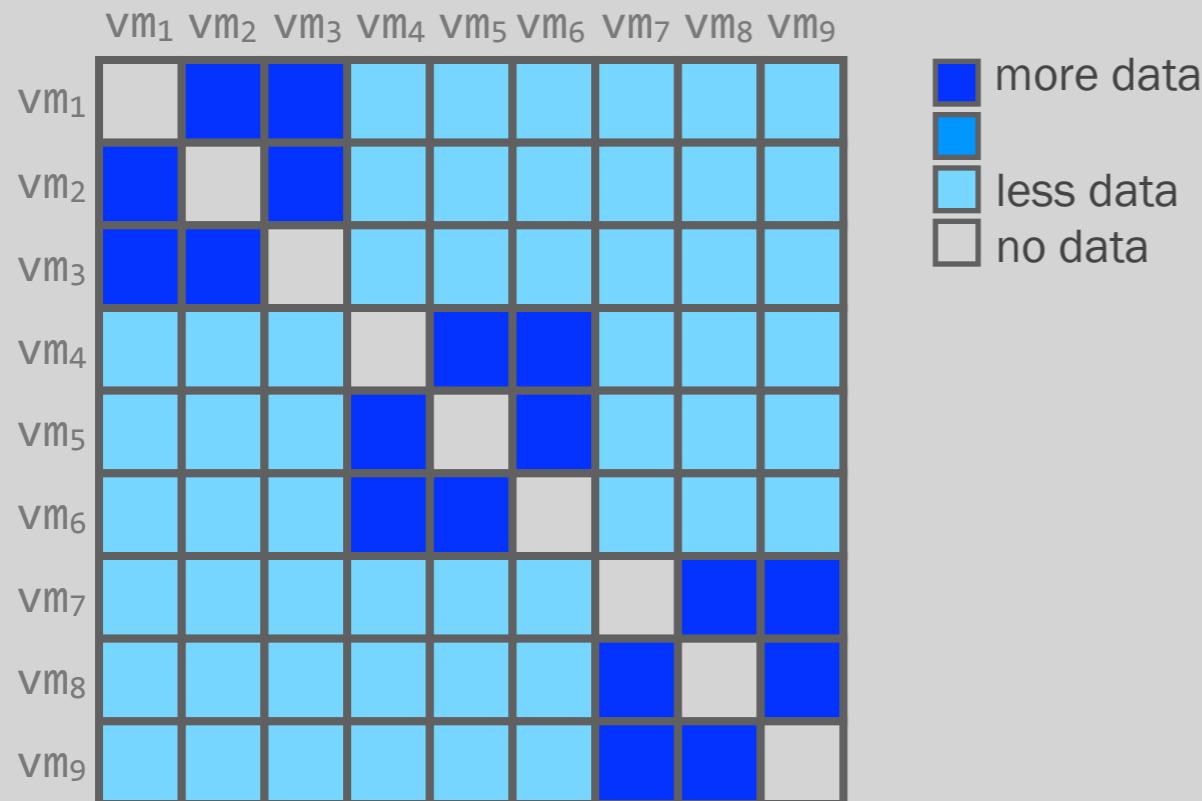
loss function
normalization

$$L(i, n) = \frac{\|\vec{M}_i - \vec{M}_n\|}{\|\vec{M}_n\|}$$

- Cicada's prediction algorithm draws inspiration from the following works:
- [1] Wang, et al. "COPE: Traffic Engineering in Dynamic Networks". SIGCOMM 2006.
 - [2] Herbster and Warmuth. "Tracking the Best Expert". Machine Learning Vol 32, 1998.

EVALUATION

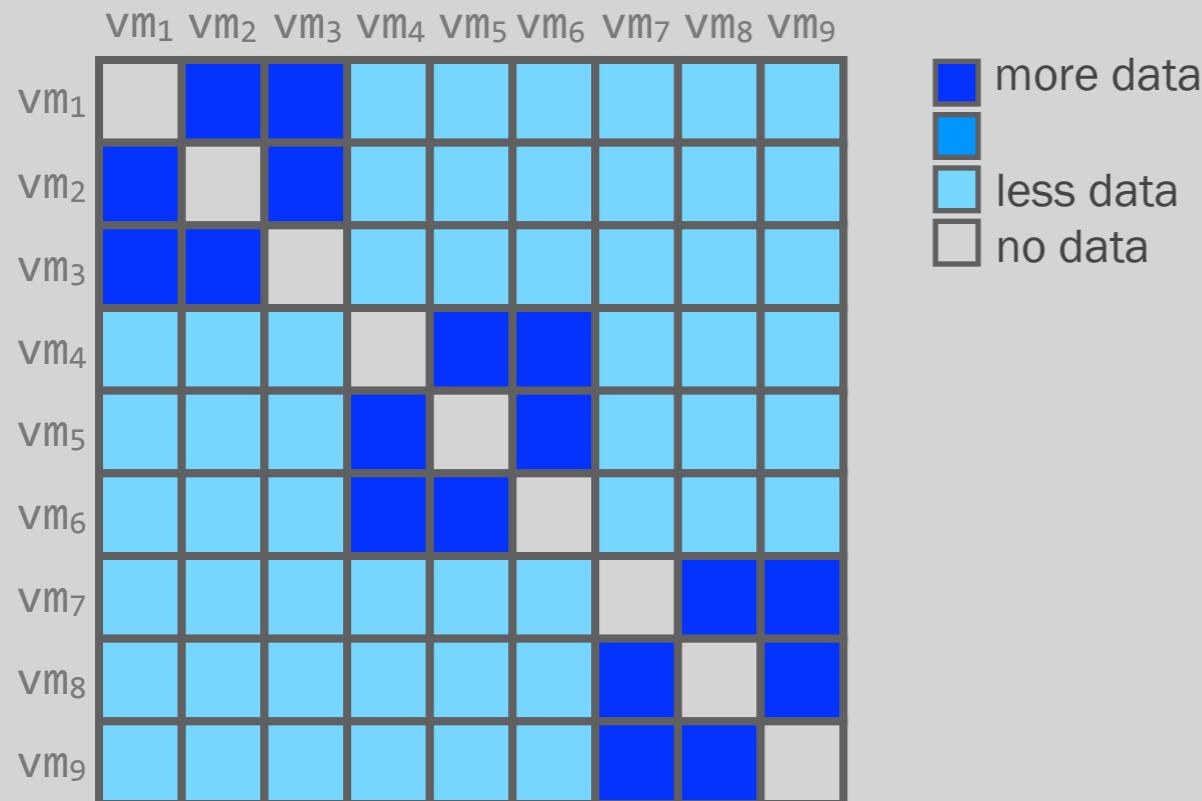
to evaluate cicada, we compared its predictions to ones made by a VOC-style system [1]



- [■] more data
- [■] less data
- [■] no data

EVALUATION

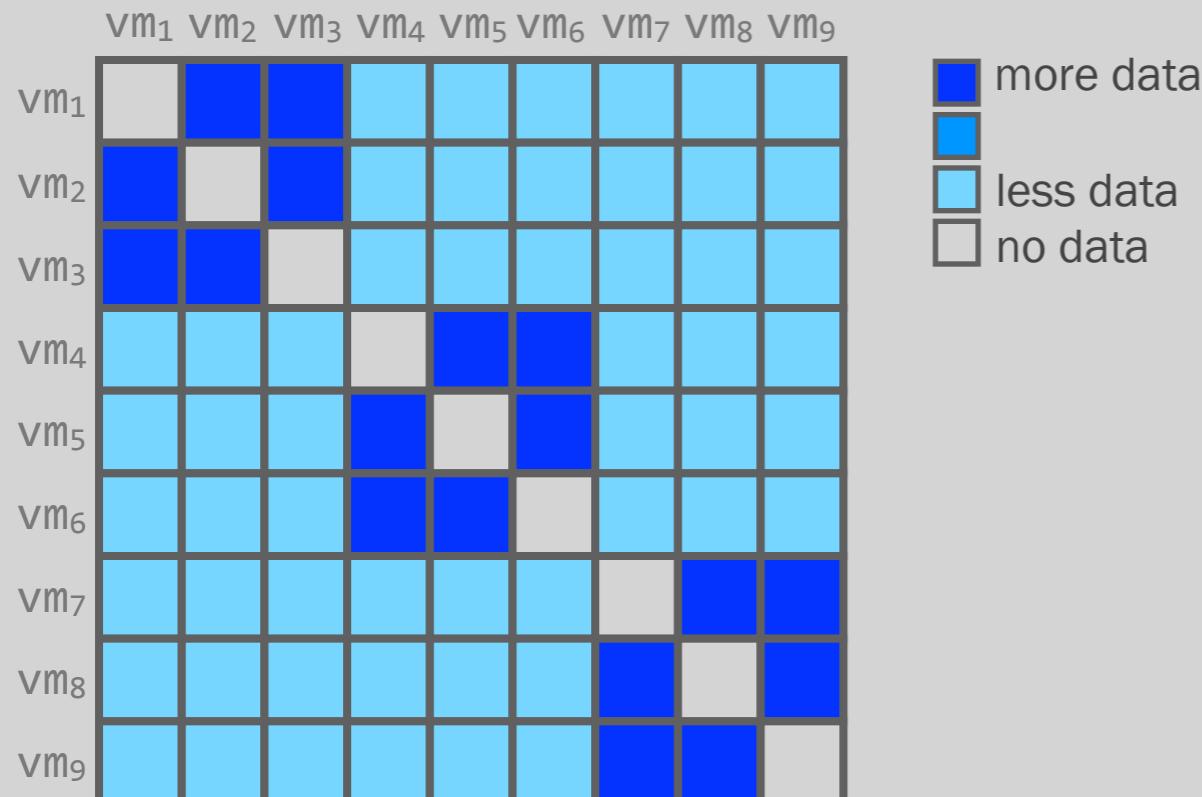
to evaluate cicada, we compared its predictions to ones made by a VOC-style system [1]



VOC requires customers to input parameters
(group sizes, amount of bandwidth)

EVALUATION

to evaluate cicada, we compared its predictions to ones made by a VOC-style system [1]

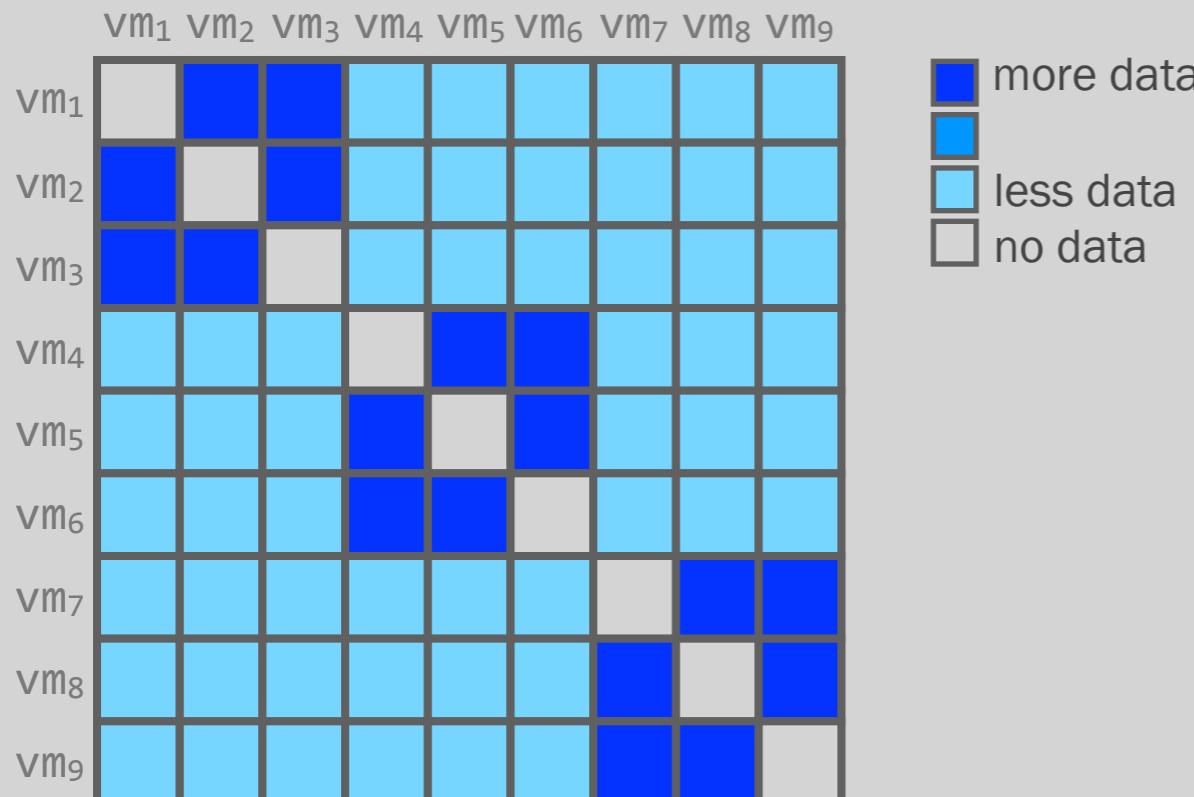


VOC requires customers to input parameters
(group sizes, amount of bandwidth)

we used an oracle method to pick the best possible parameters, and also allowed for temporal variability

EVALUATION

to evaluate cicada, we compared its predictions to ones made by a VOC-style system [1]



VOC requires customers to input parameters

(group sizes, amount of bandwidth)

we used an oracle method to pick the best possible parameters, and also allowed for temporal variability

L2-norm error

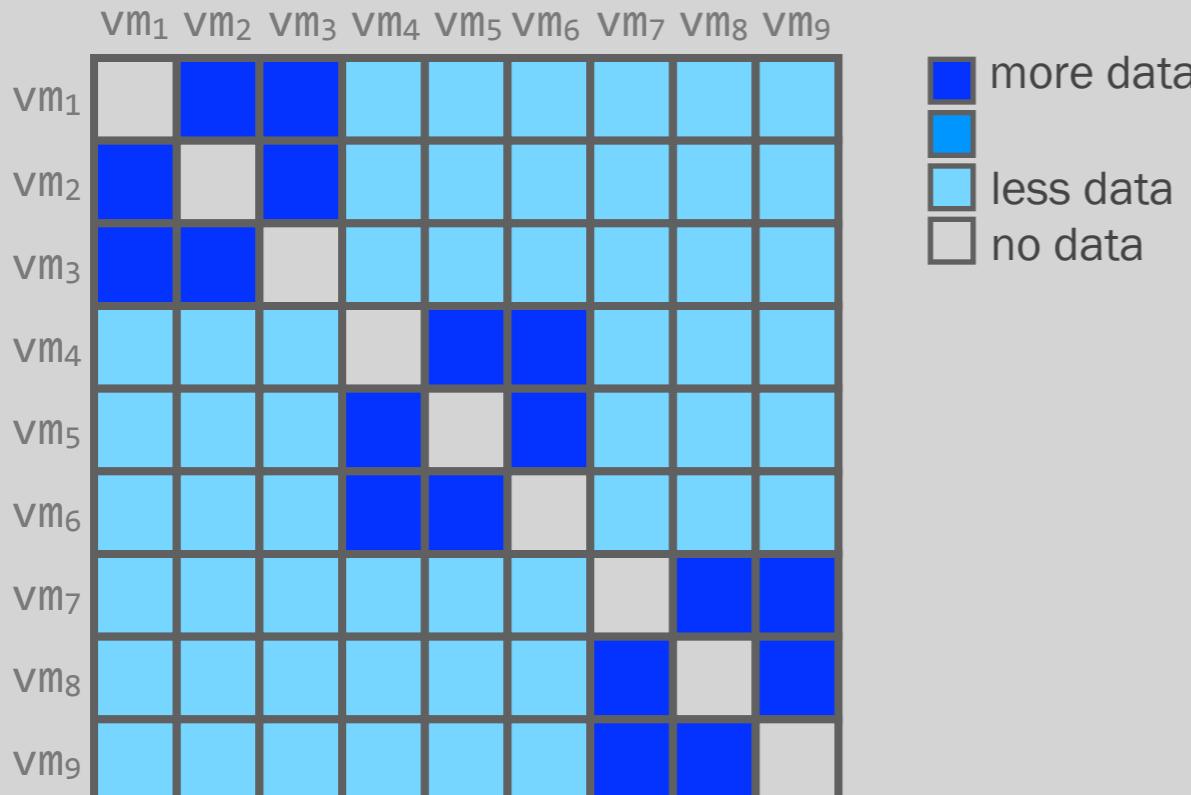
$$E = \frac{\|\hat{M} - M\|}{\|\hat{M}\|}$$

relative error

$$E = \frac{\sum_{i=1}^{n^2} \hat{M}[i] - M[i]}{\sum_{i=1}^{n^2} M[i]}$$

EVALUATION

to evaluate cicada, we compared its predictions to ones made by a VOC-style system [1]



VOC requires customers to input parameters
(group sizes, amount of bandwidth)

we used an oracle method to pick the best possible parameters, and also allowed for temporal variability

L2-norm error

$$E = \frac{\|\hat{M} - M\|}{\|\hat{M}\|}$$

relative error

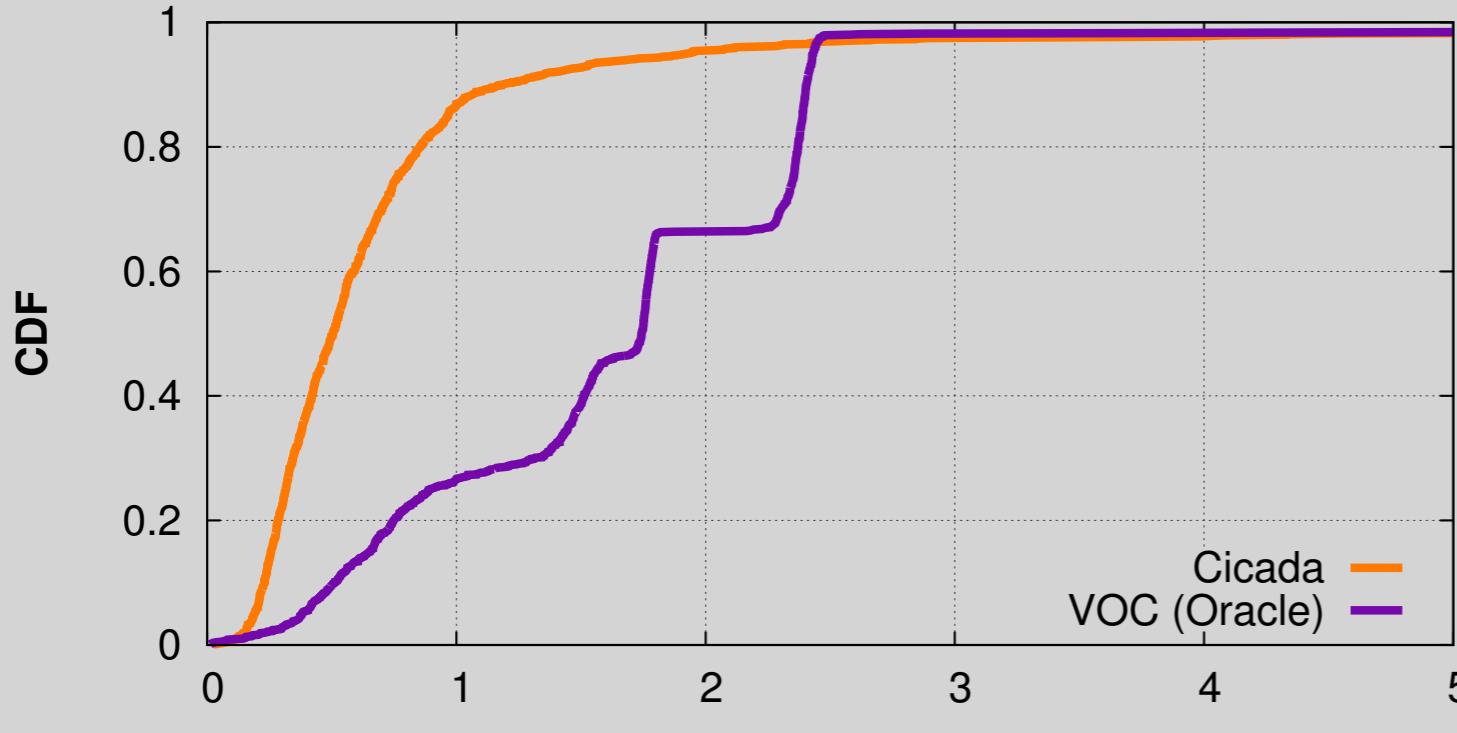
$$E = \frac{\sum_{i=1}^{n^2} \hat{M}[i] - M[i]}{\sum_{i=1}^{n^2} M[i]}$$



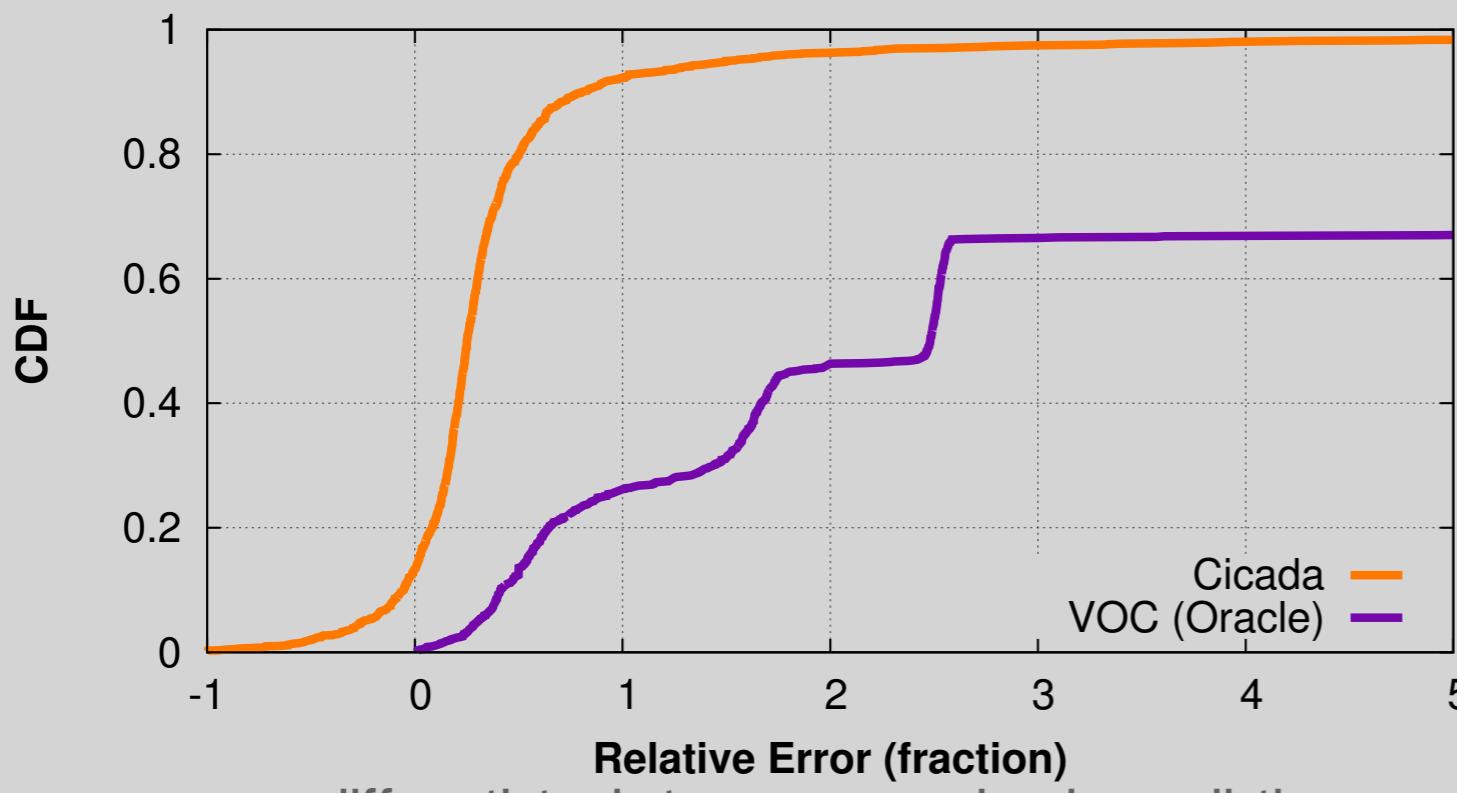
differentiates between over- and under-prediction

RESULTS

PREDICTING AVERAGE DEMAND



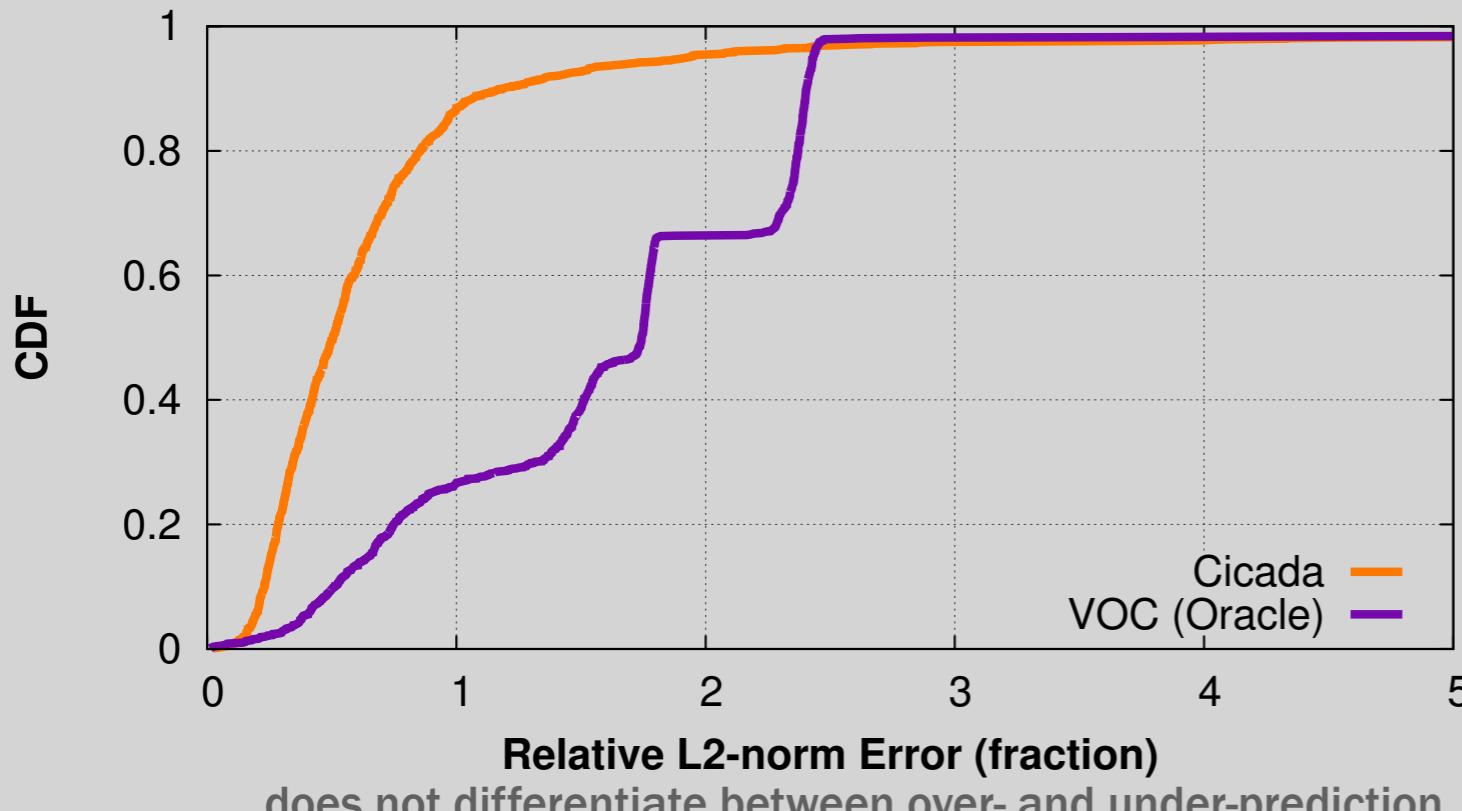
does not differentiate between over- and under-prediction



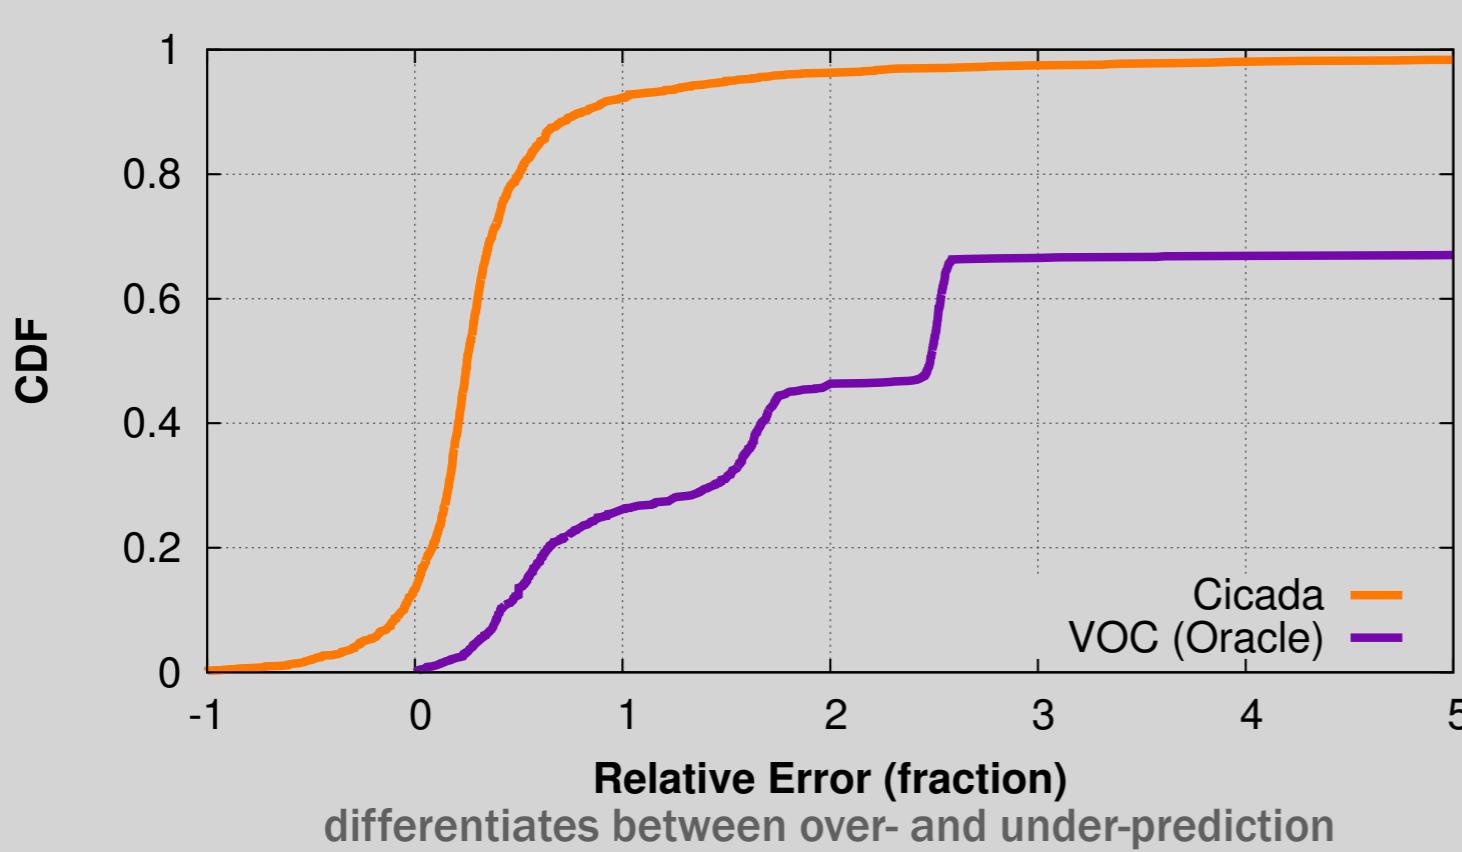
differentiates between over- and under-prediction

RESULTS

PREDICTING AVERAGE DEMAND



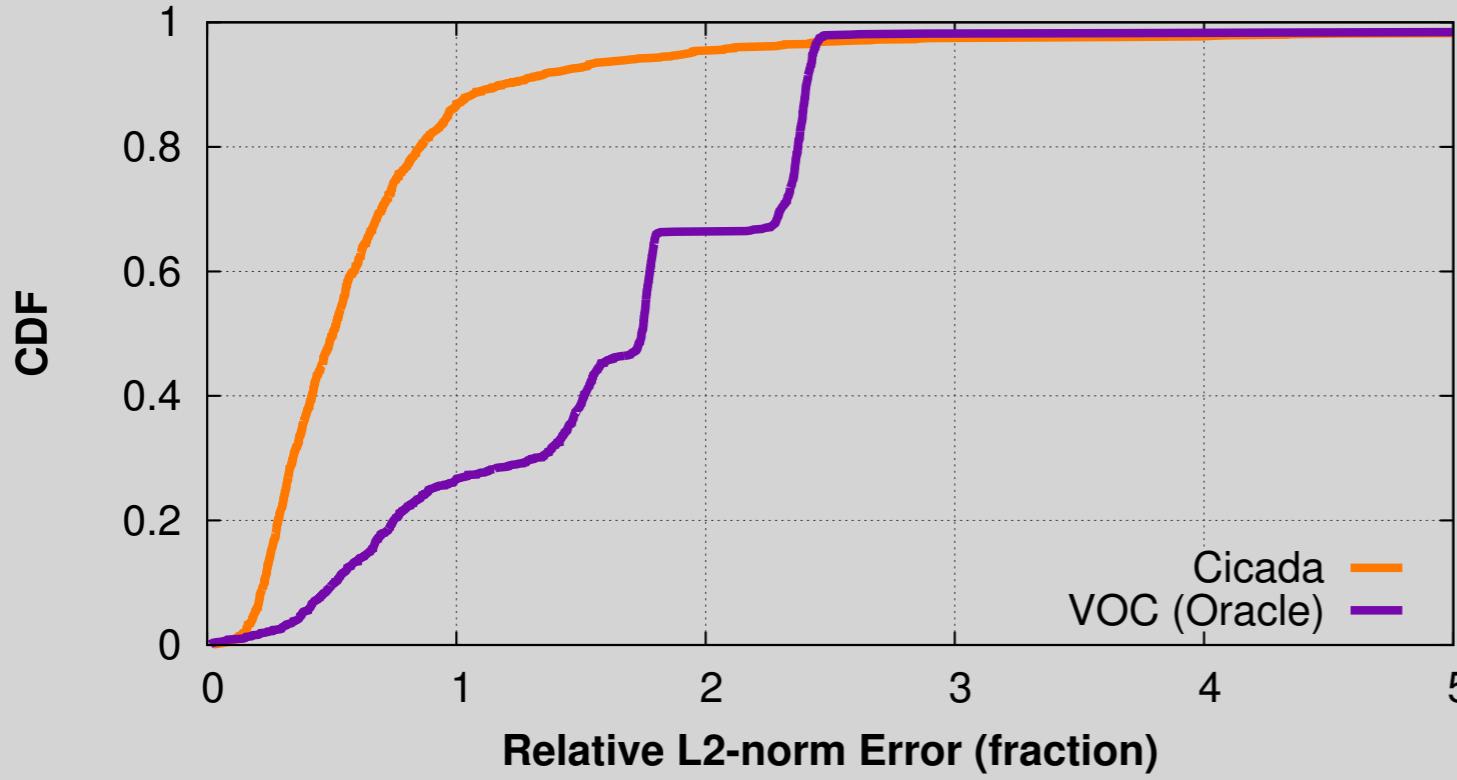
cicada's predictions outperform VOC-style predictions (median error decreases by 90%) and require no customer input
(71% for L2 error)



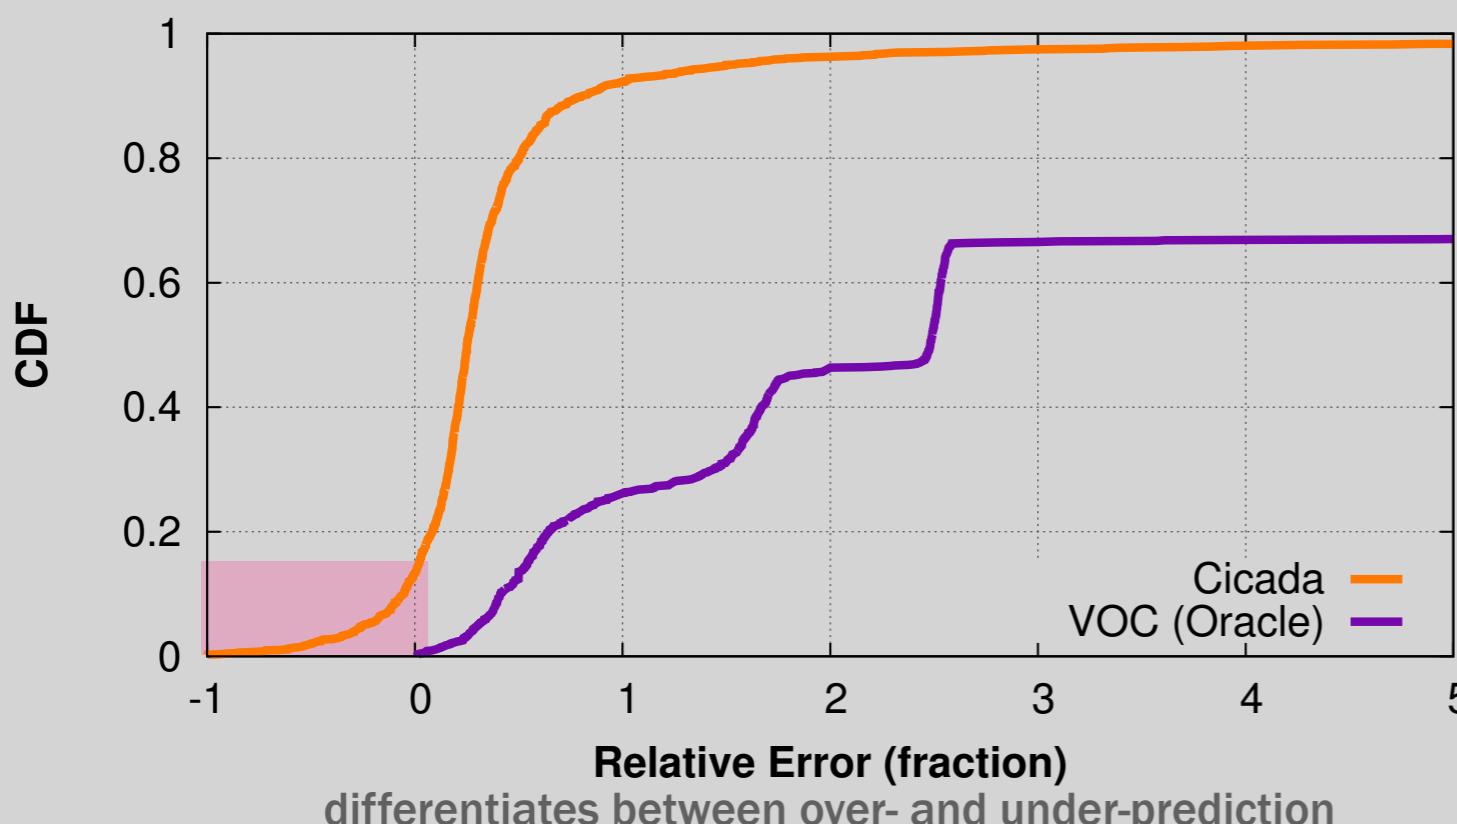
differentiates between over- and under-prediction

RESULTS

PREDICTING AVERAGE DEMAND



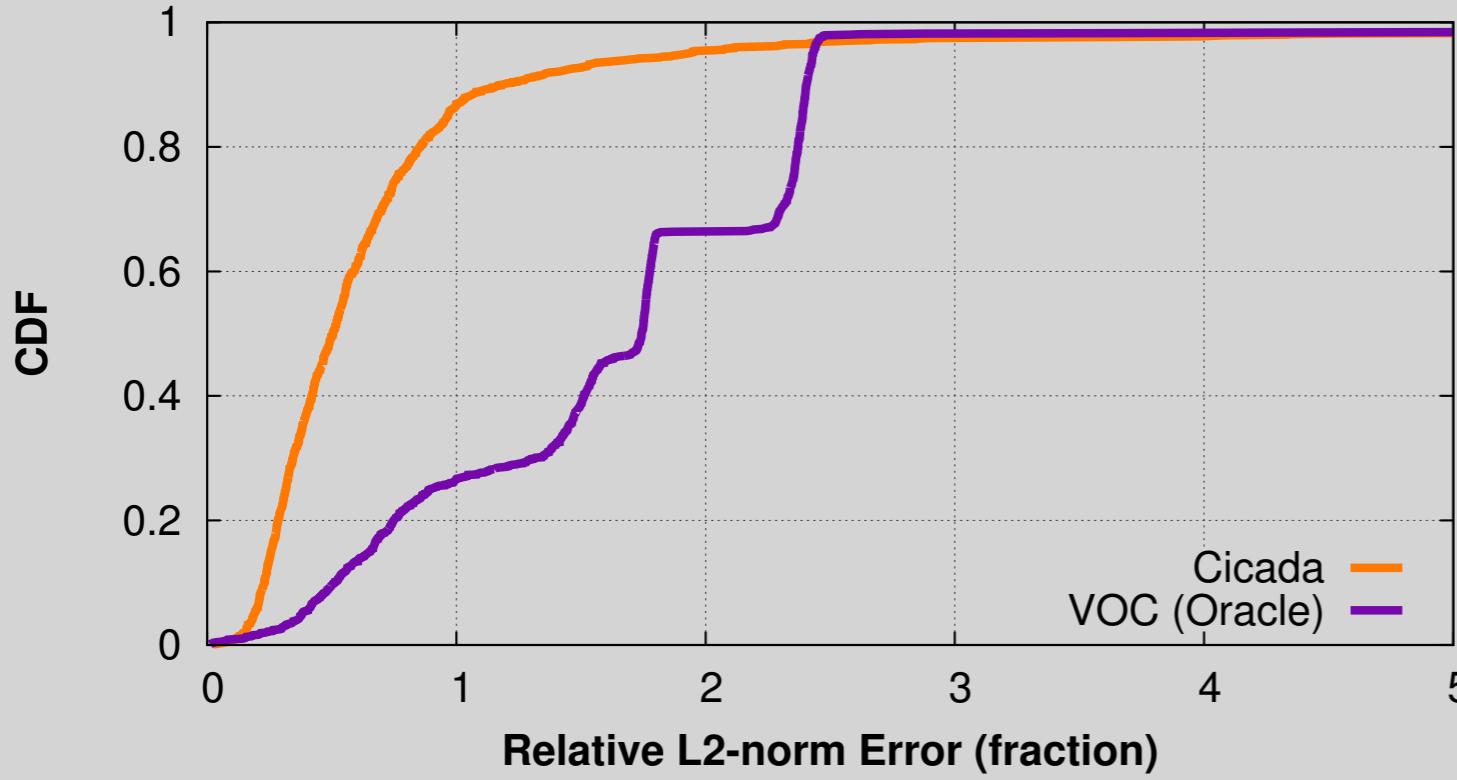
cicada's predictions outperform VOC-style predictions (median error decreases by 90%) and require no customer input
(71% for L2 error)



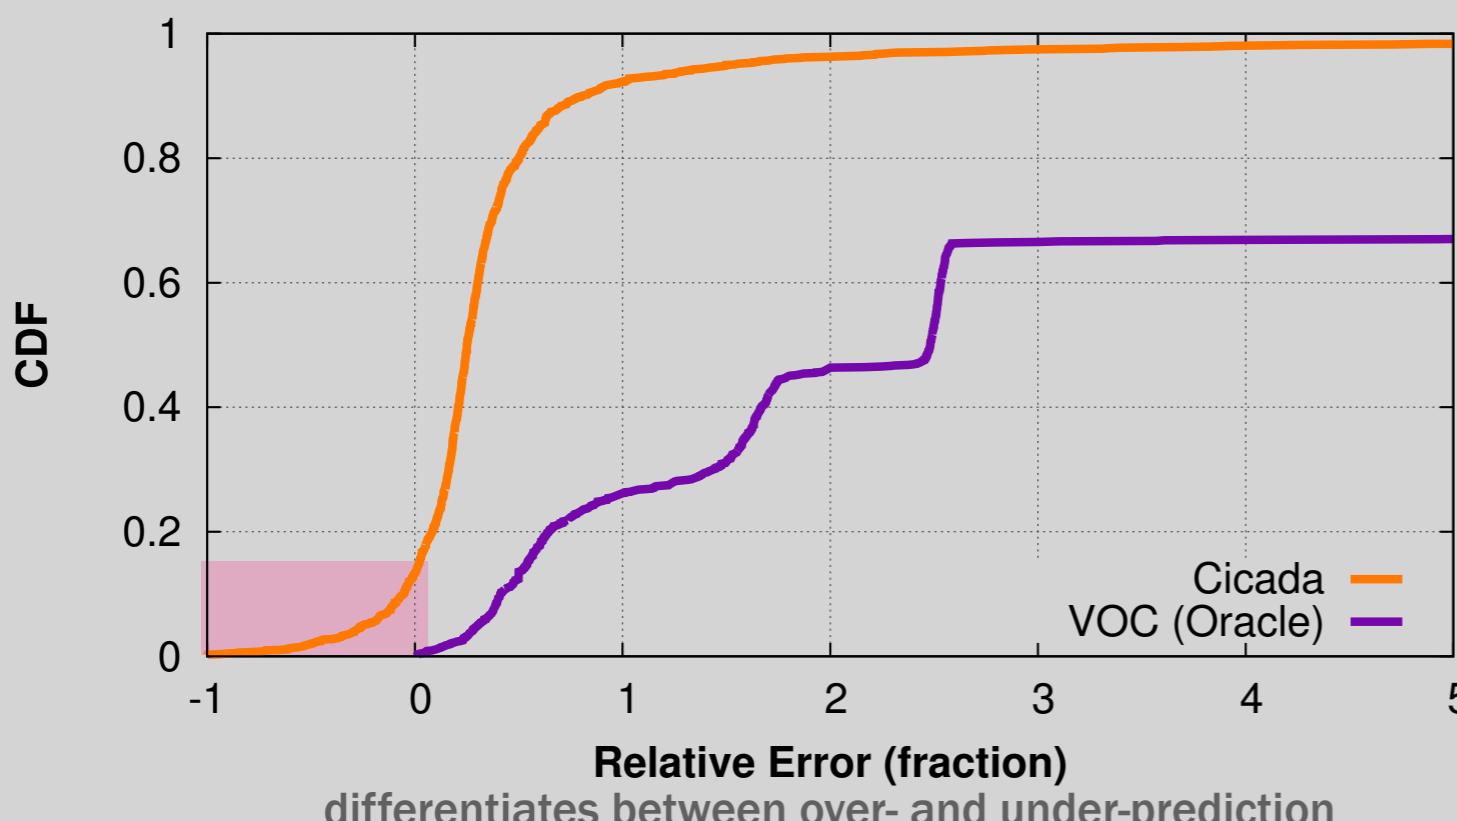
cicada occasionally under-predicts; VOC does not because we selected its parameters to avoid under-prediction

RESULTS

PREDICTING AVERAGE DEMAND



cicada's predictions outperform VOC-style predictions (median error decreases by 90%) and require no customer input
(71% for L2 error)



cicada occasionally under-predicts; VOC does not because we selected its parameters to avoid under-prediction

cicada's predictions typically require only 1-2 hours of application history

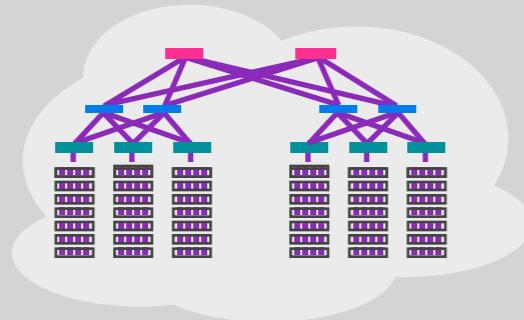
PREDICTIONS → GUARANTEES

cicada turns its predictions into **pipe-model** guarantees;
a separate guarantee for each (source, destination) pair

PREDICTIONS → GUARANTEES

cicada turns its predictions into **pipe-model** guarantees;
a separate guarantee for each (source, destination) pair

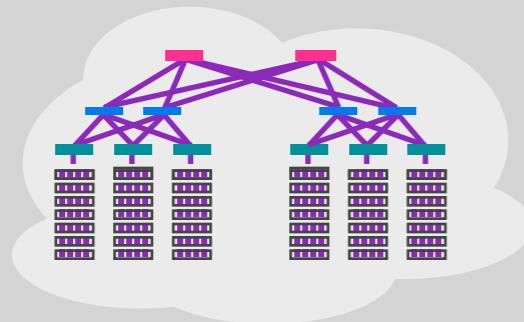
network resources must be available



PREDICTIONS → GUARANTEES

cicada turns its predictions into **pipe-model** guarantees;
a separate guarantee for each (source, destination) pair

network resources must be available



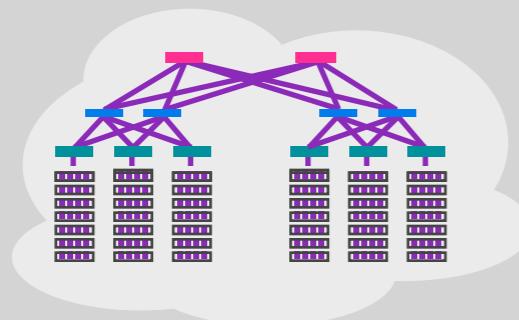
customers can add a buffer to
the offered guarantees



PREDICTIONS → GUARANTEES

cicada turns its predictions into **pipe-model** guarantees;
a separate guarantee for each (source, destination) pair

network resources must be available



customers can add a buffer to
the offered guarantees



the provider may choose to augment
predictions that cicada is not confident about

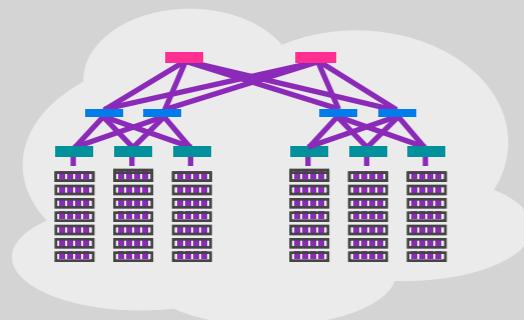


typically not confident for “small” applications

PREDICTIONS → GUARANTEES

cicada turns its predictions into **pipe-model** guarantees;
a separate guarantee for each (source, destination) pair

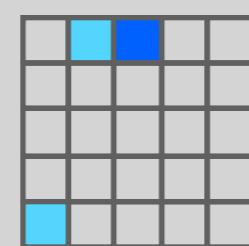
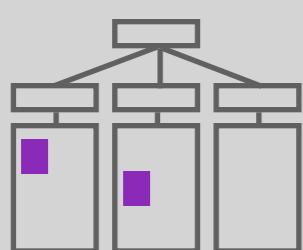
network resources must be available



customers can add a buffer to
the offered guarantees

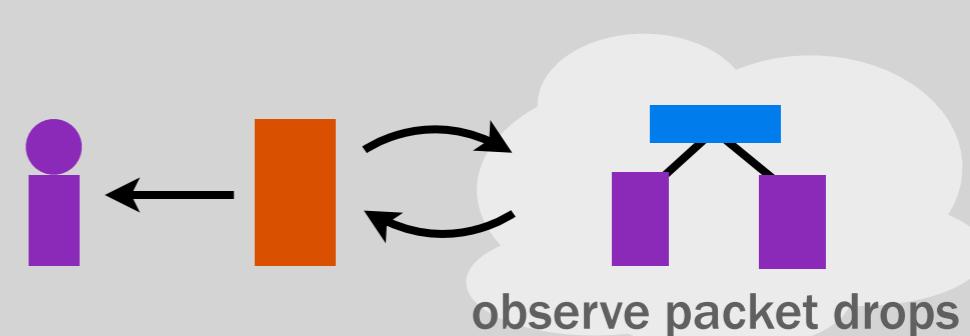


the provider may choose to augment
predictions that cicada is not confident about



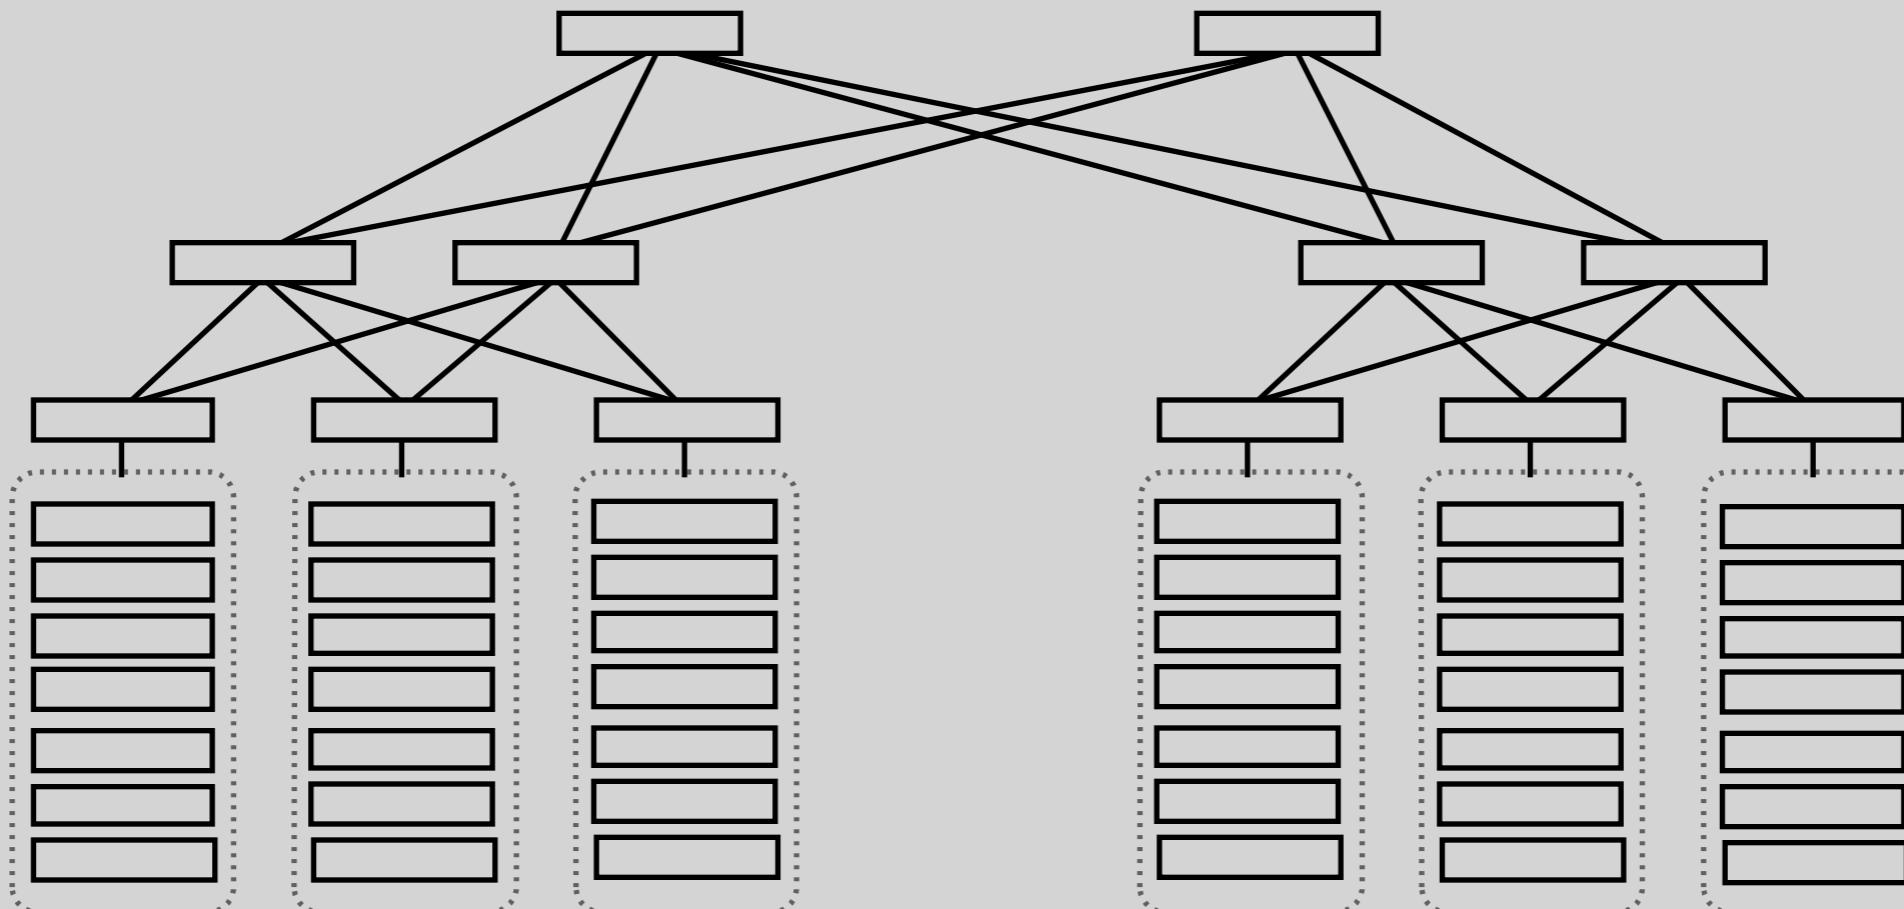
typically not confident for “small” applications

cicada can detect when a guarantee
is too low, and offer a new one



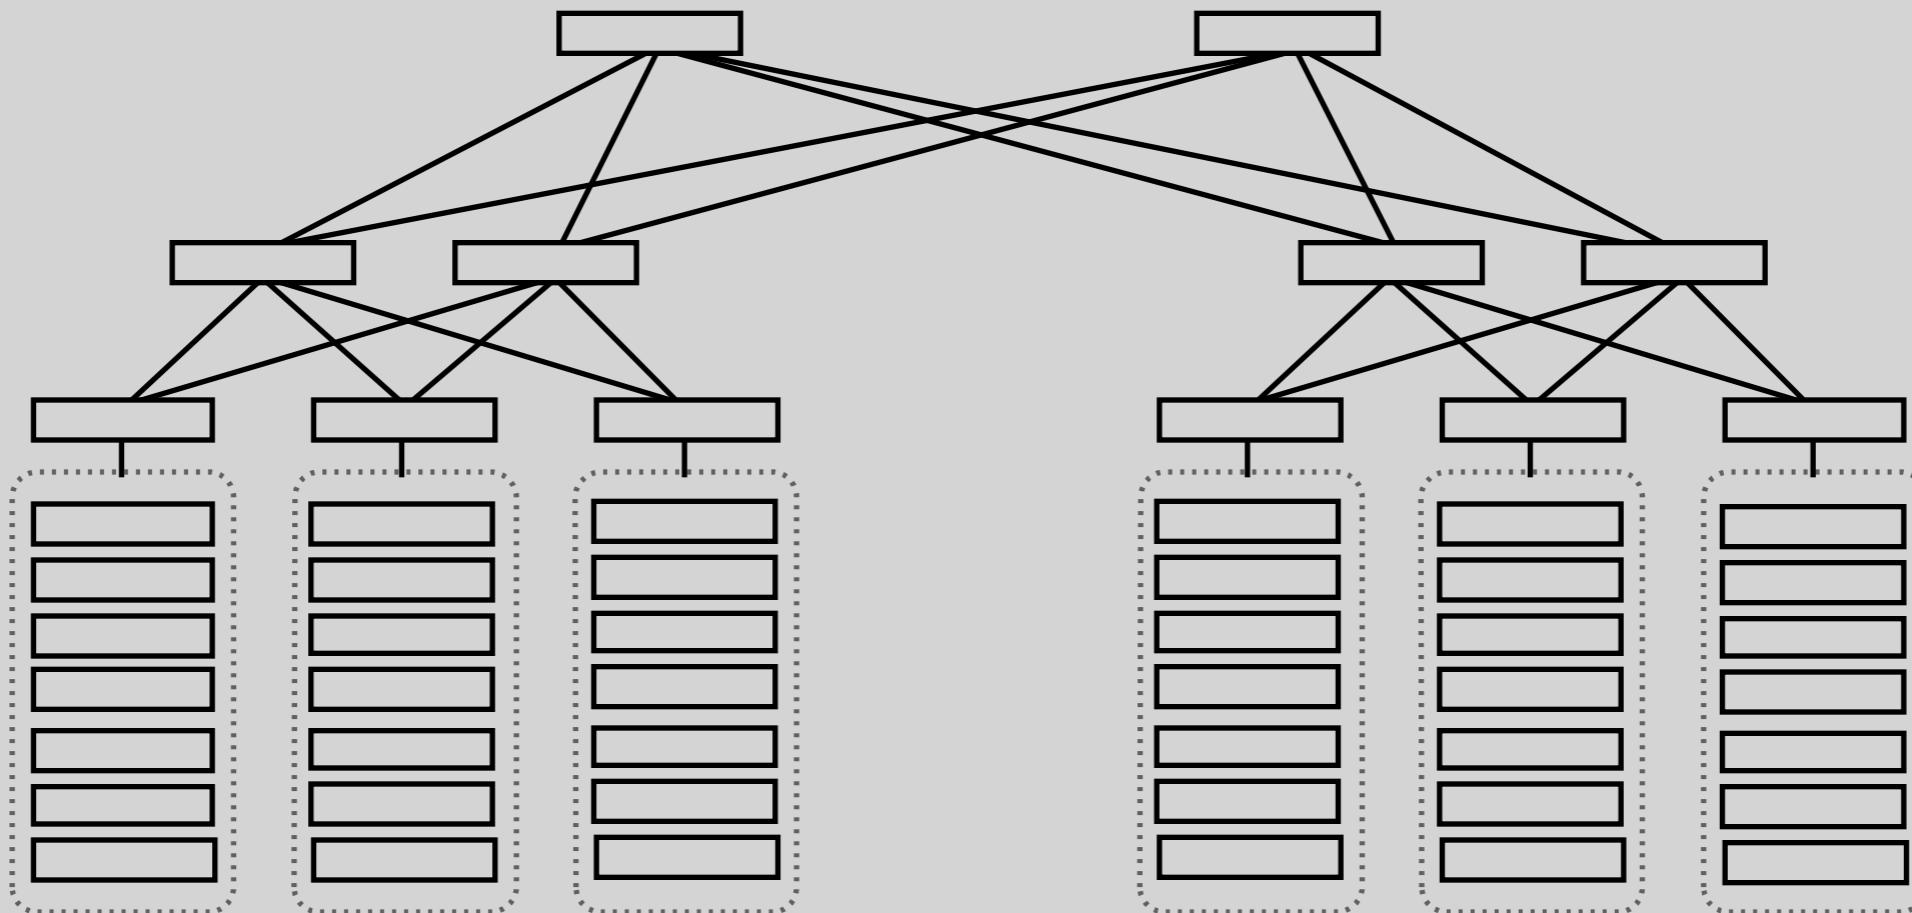
NETWORK UTILIZATION

do cicada's bandwidth guarantees improve network utilization?



NETWORK UTILIZATION

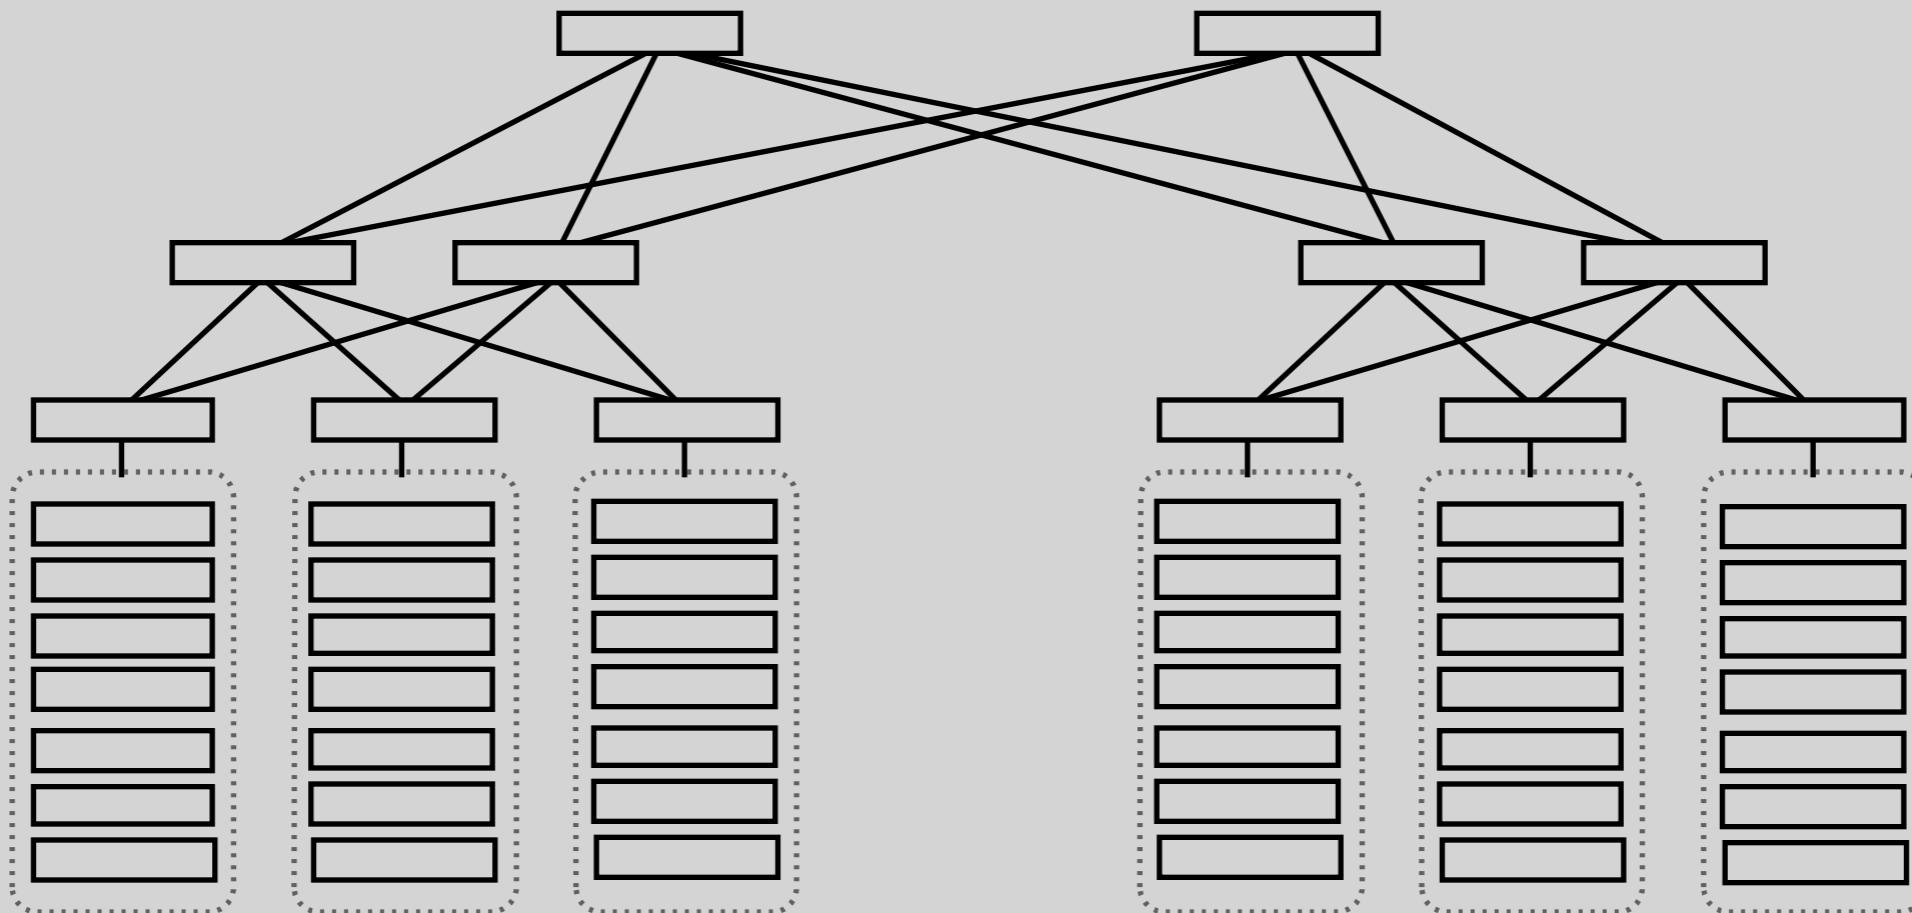
do cicada's bandwidth guarantees improve network utilization?



wasted bandwidth: bandwidth that is guaranteed for an application, but not *used* by the application

NETWORK UTILIZATION

do cicada's bandwidth guarantees improve network utilization?

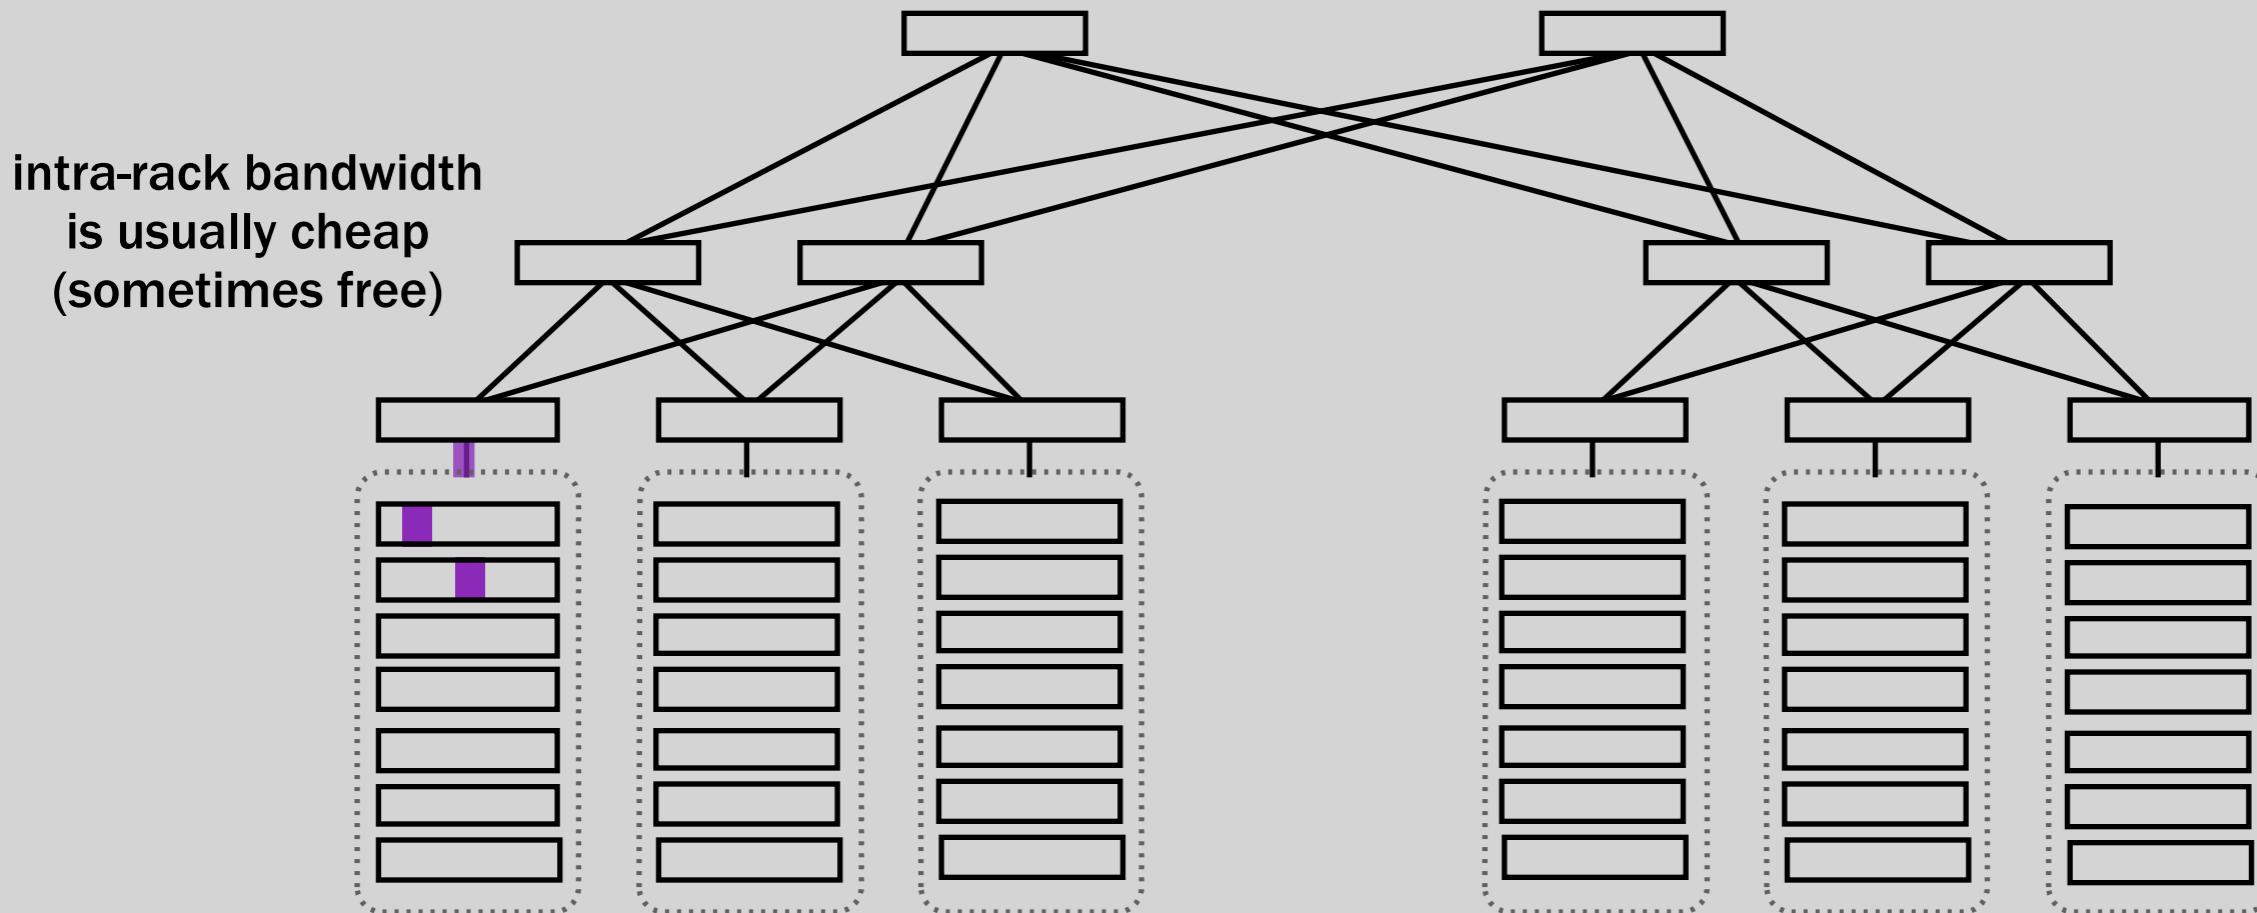


wasted bandwidth: bandwidth that is guaranteed for an application, but not *used* by the application

provider's goal: minimize wasted bandwidth

NETWORK UTILIZATION

do cicada's bandwidth guarantees improve network utilization?

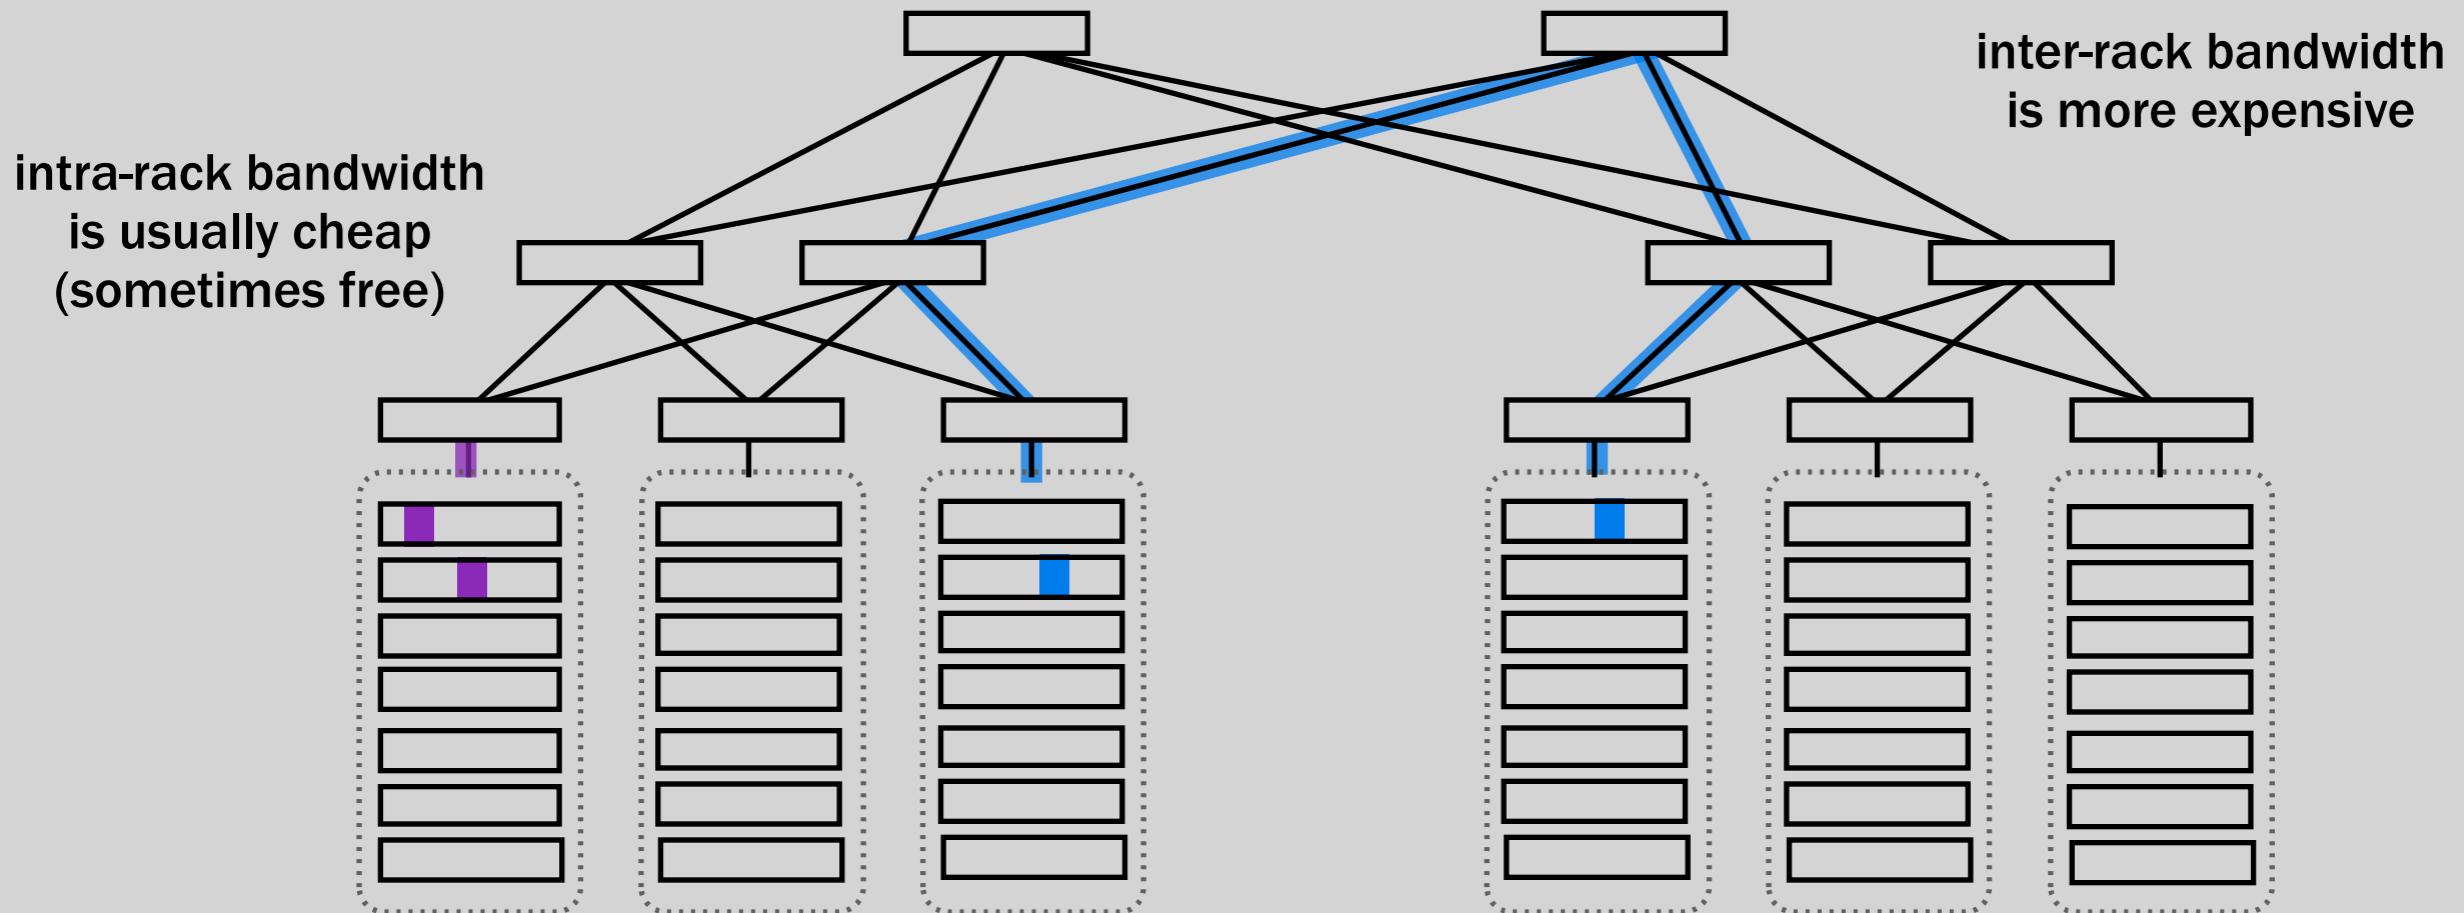


wasted bandwidth: bandwidth that is guaranteed for an application, but not *used* by the application

provider's goal: minimize wasted bandwidth

NETWORK UTILIZATION

do cicada's bandwidth guarantees improve network utilization?

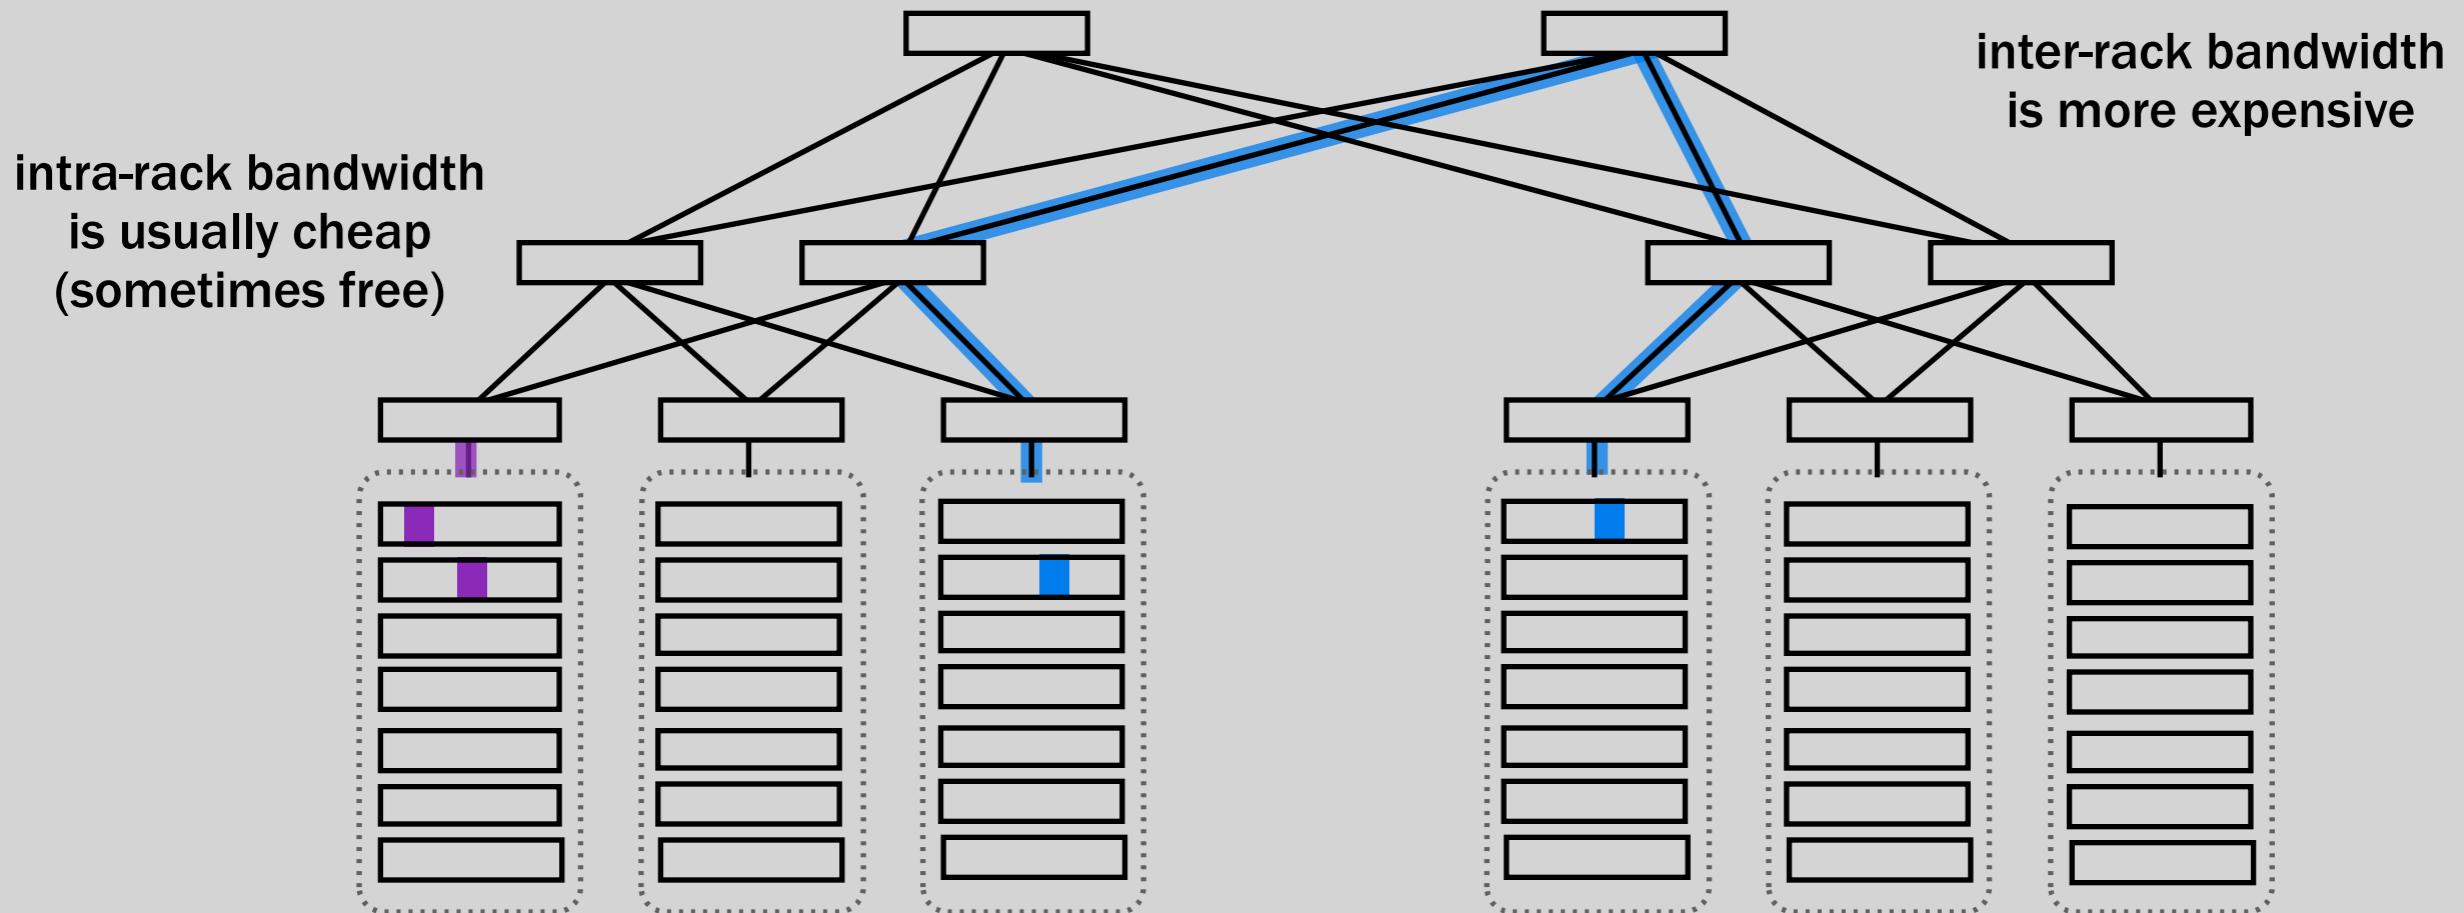


wasted bandwidth: bandwidth that is guaranteed for an application, but not *used* by the application

provider's goal: minimize wasted bandwidth

NETWORK UTILIZATION

do cicada's bandwidth guarantees improve network utilization?

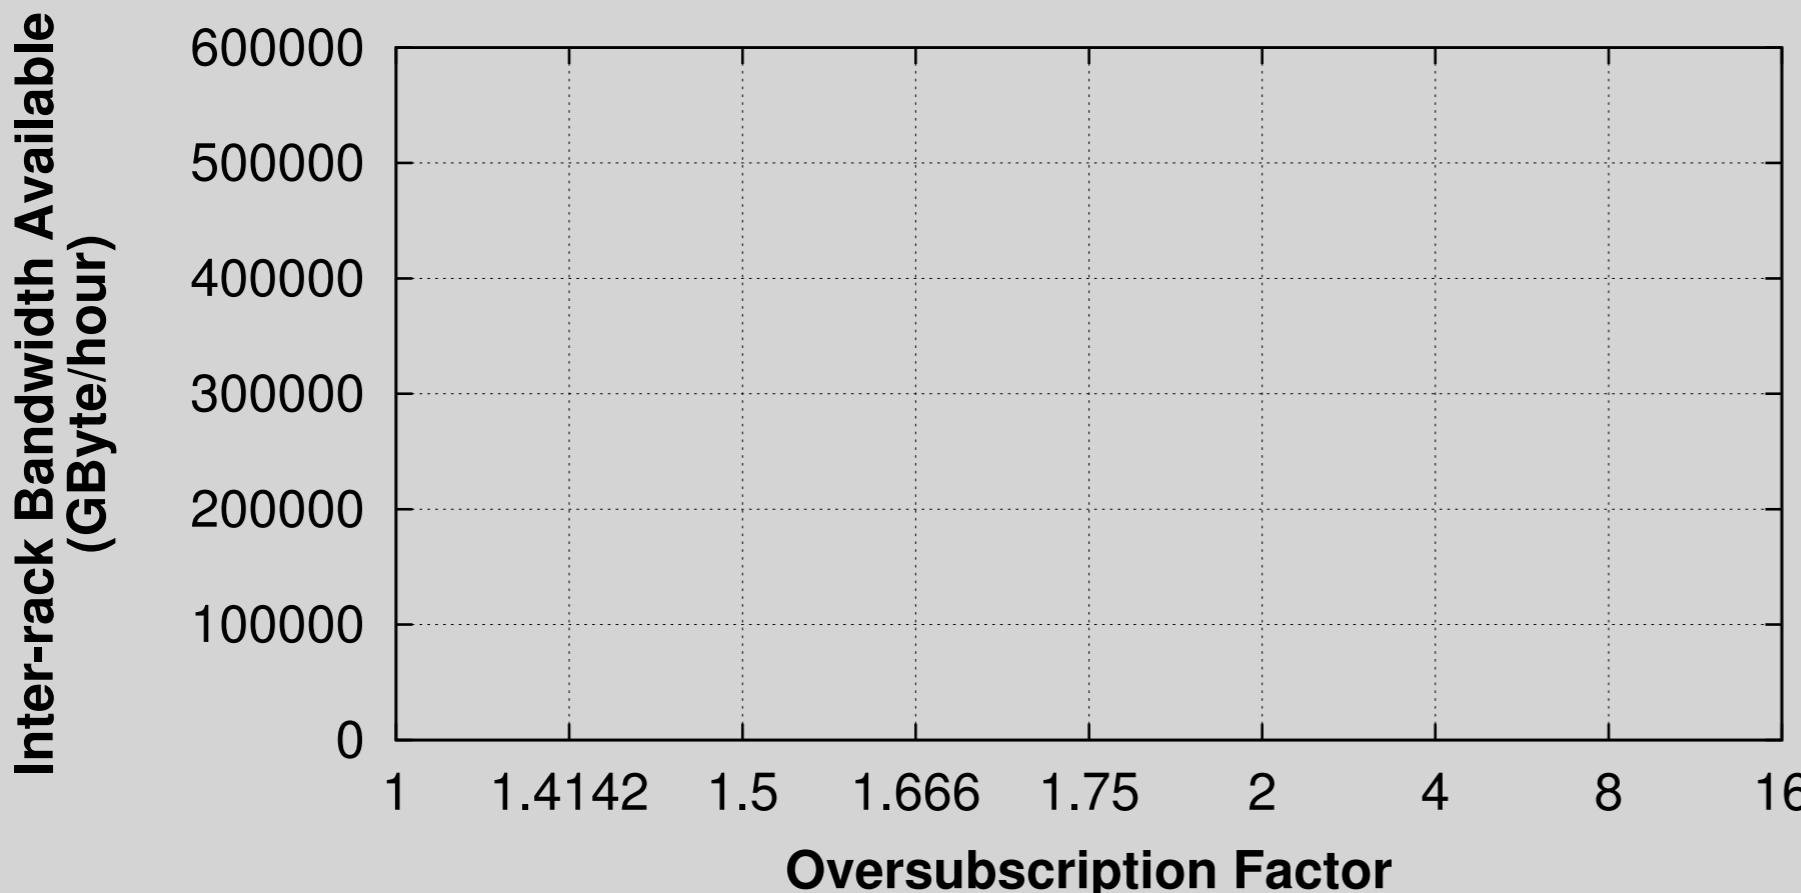


wasted bandwidth: bandwidth that is guaranteed for an application, but not *used* by the application

provider's goal: minimize wasted *inter-rack* bandwidth

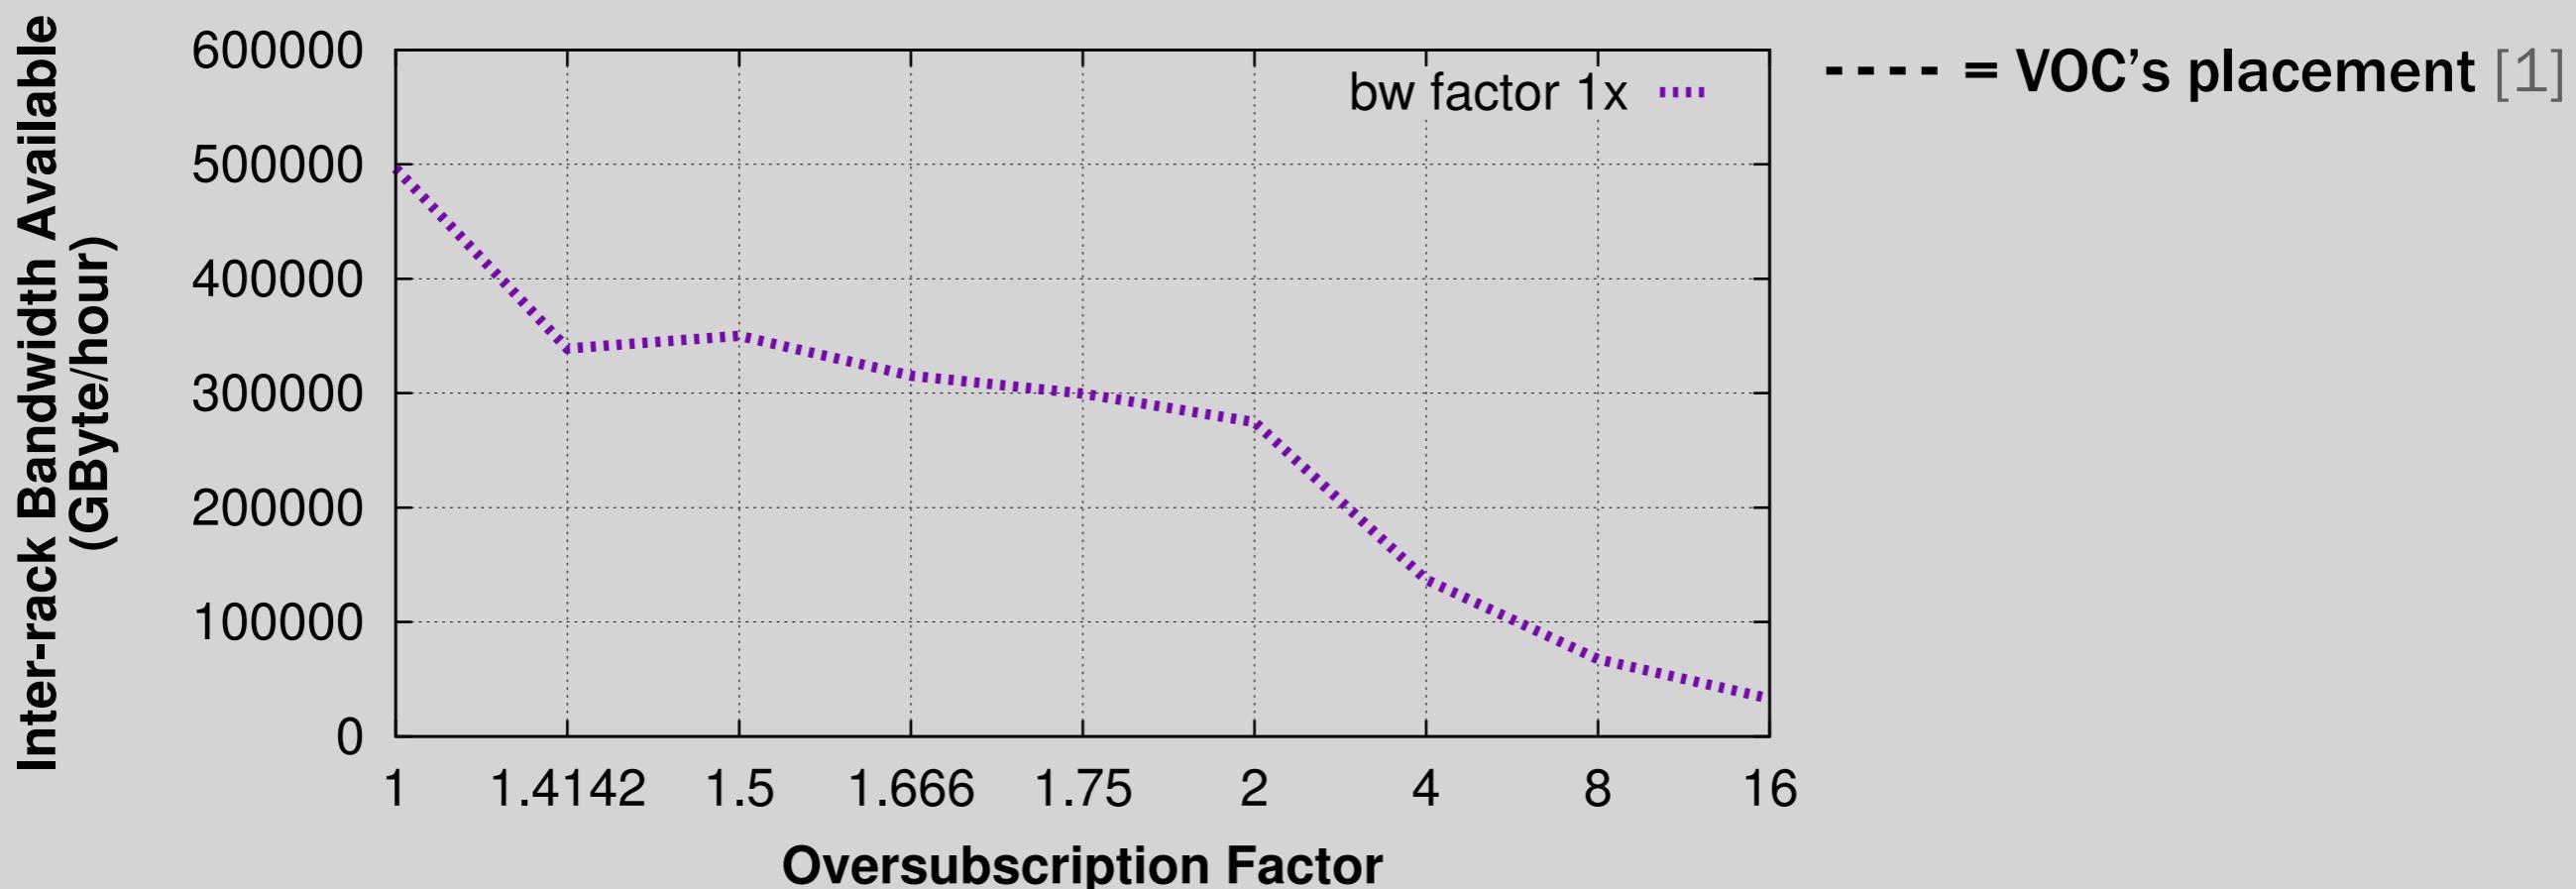
NETWORK UTILIZATION

cicada's **greedy** placement heuristic: place the pairs of VMs that need the highest guarantees on the fastest paths (see paper for details)



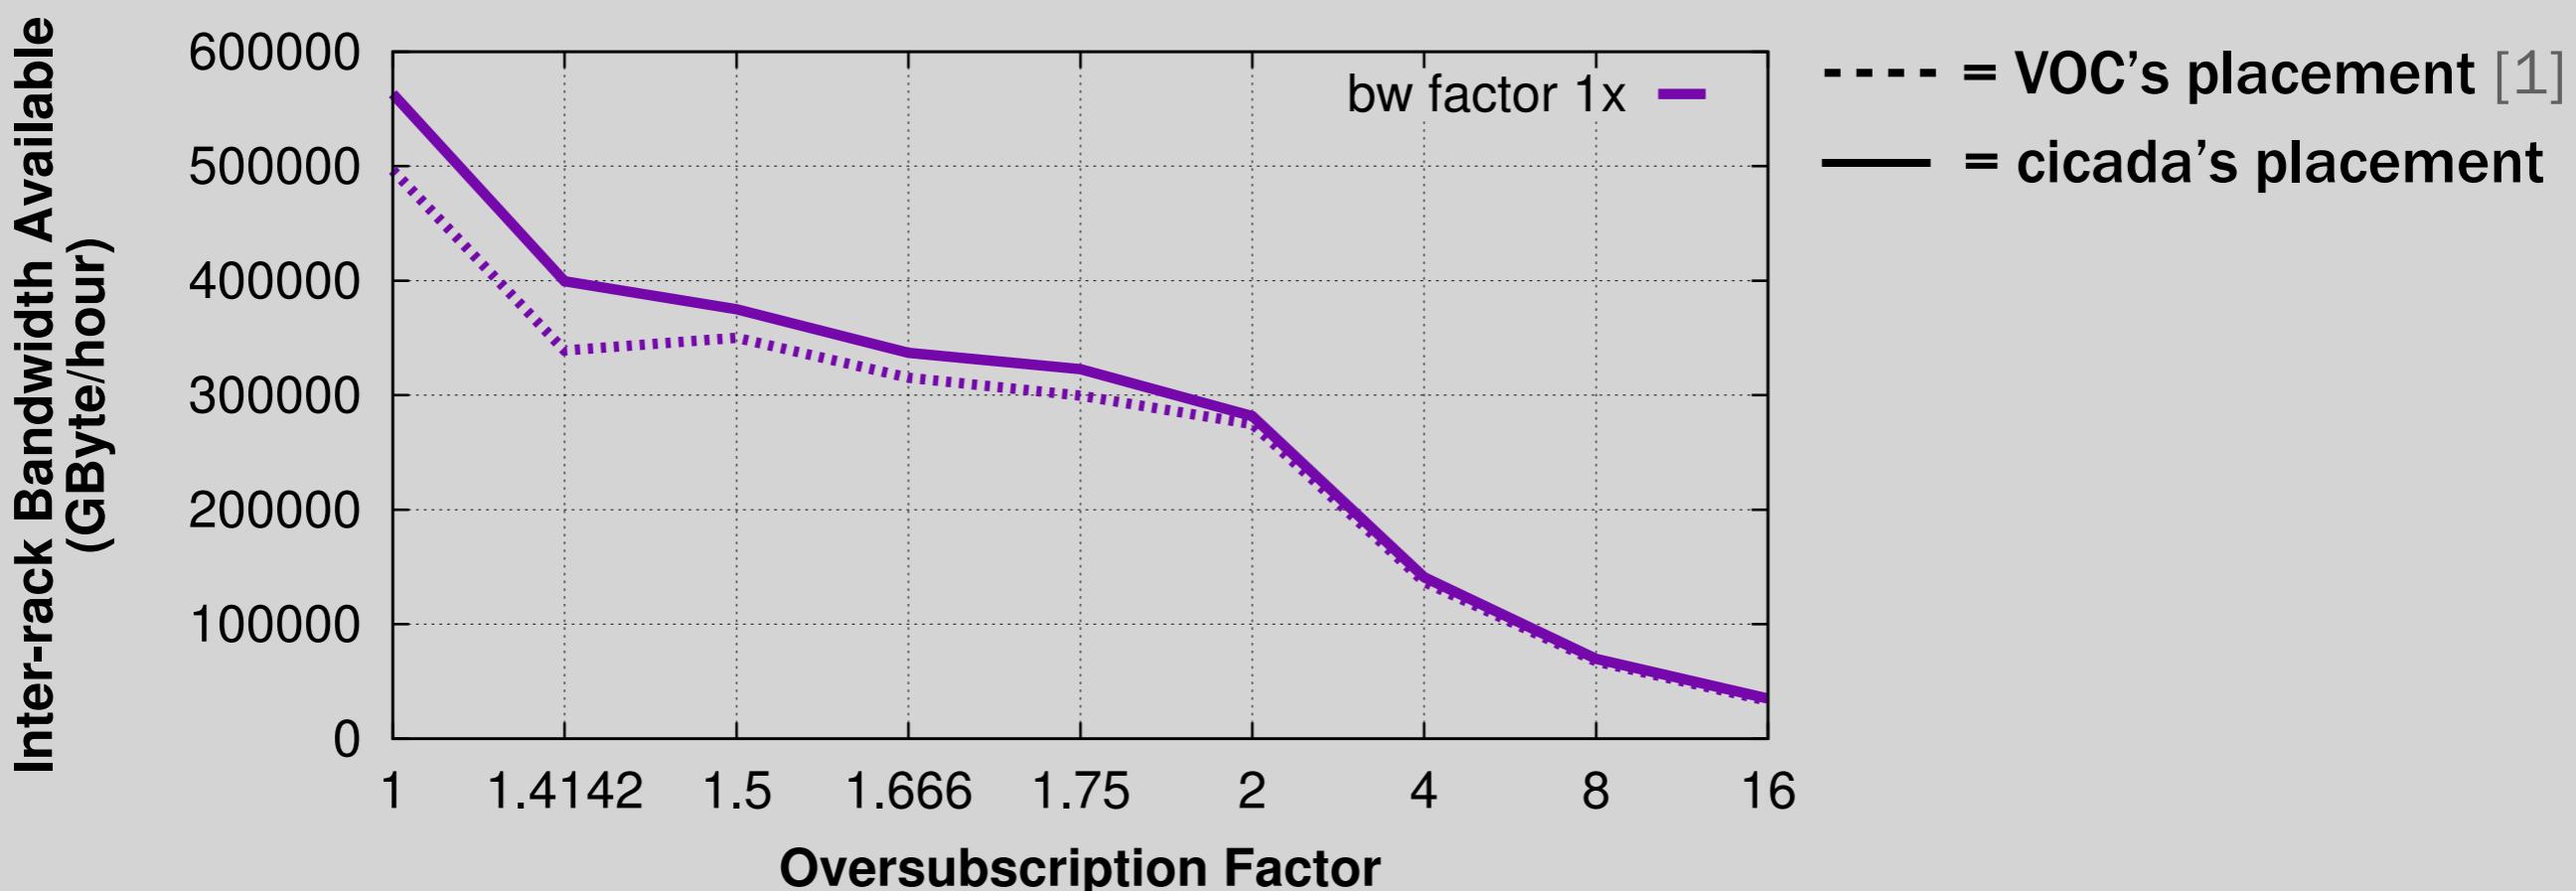
NETWORK UTILIZATION

cicada's **greedy** placement heuristic: place the pairs of VMs that need the highest guarantees on the fastest paths (see paper for details)



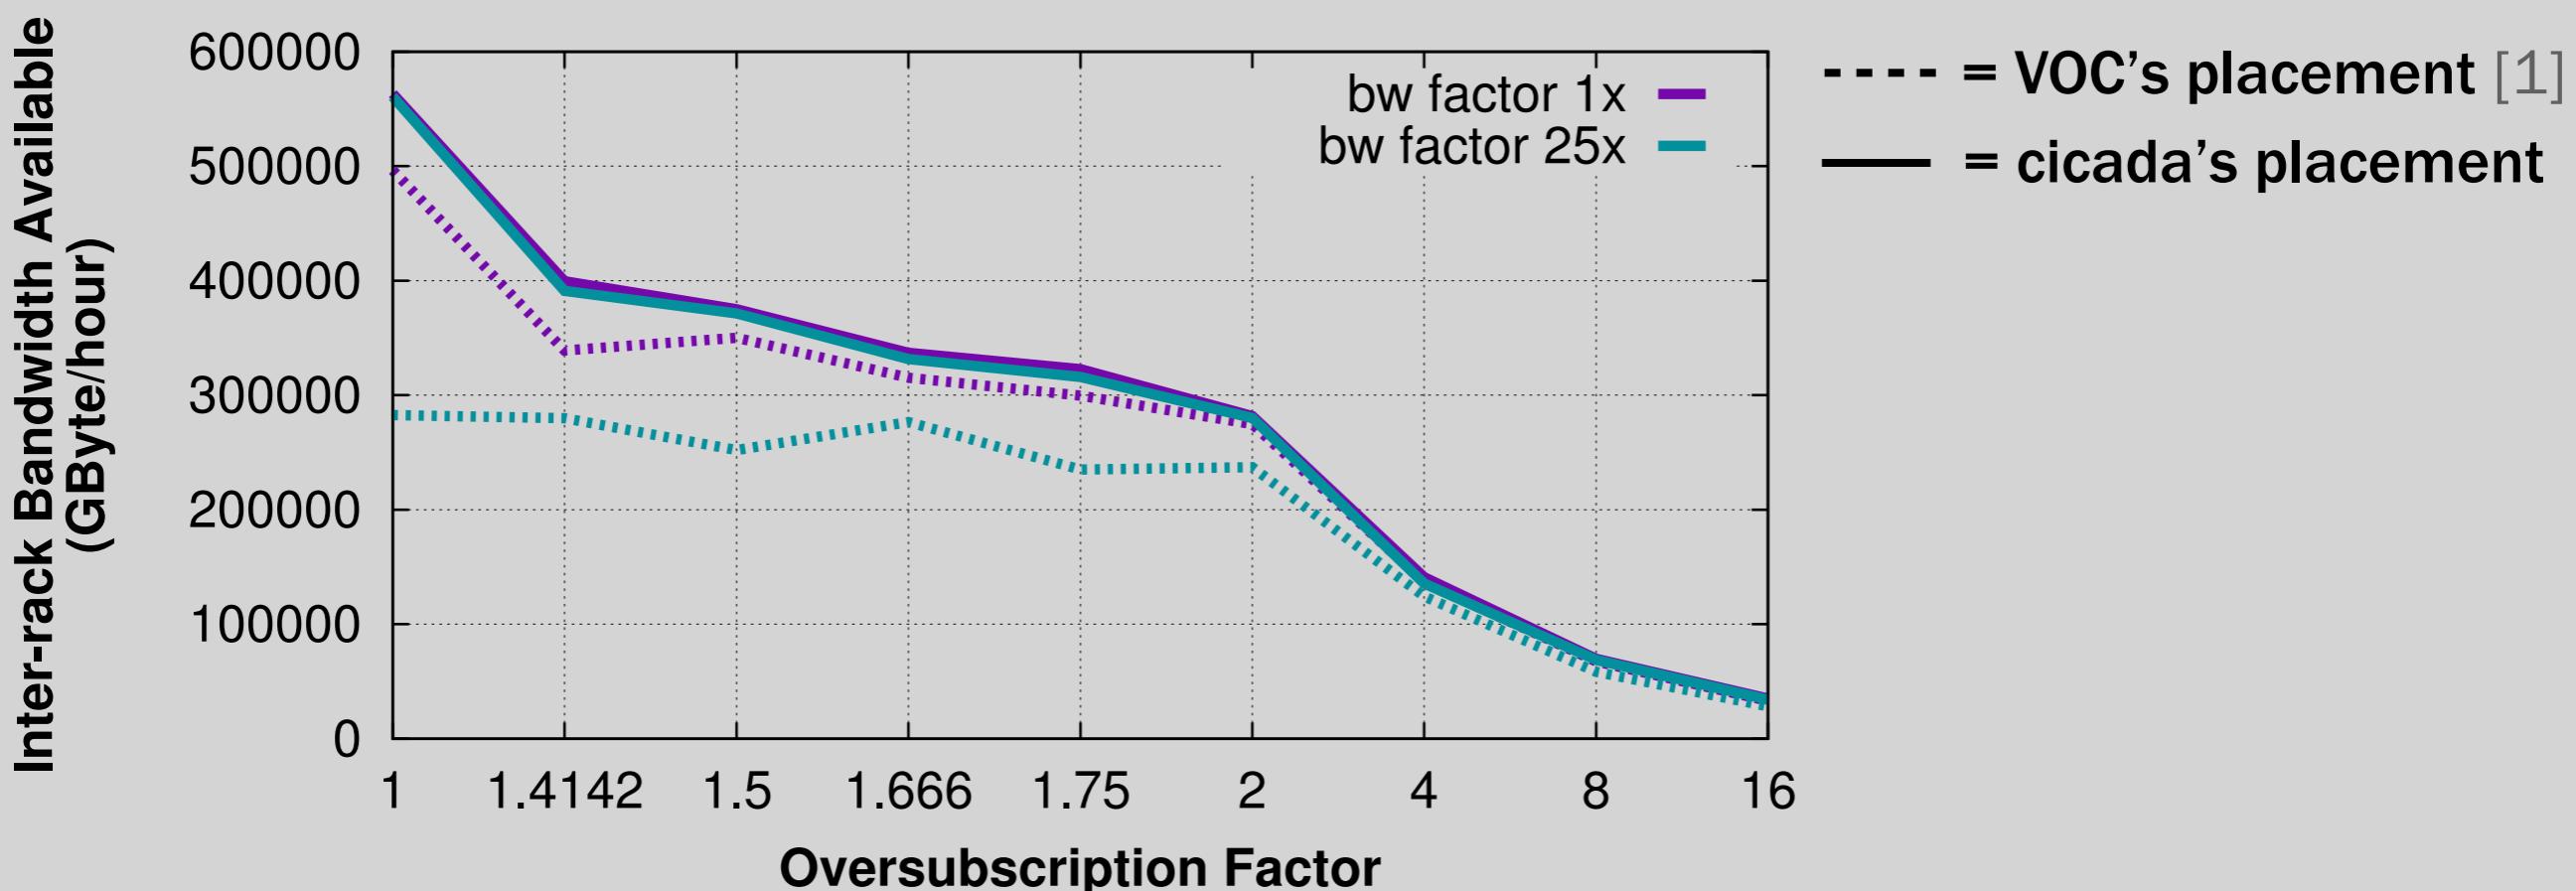
NETWORK UTILIZATION

cicada's **greedy** placement heuristic: place the pairs of VMs that need the highest guarantees on the fastest paths (see paper for details)



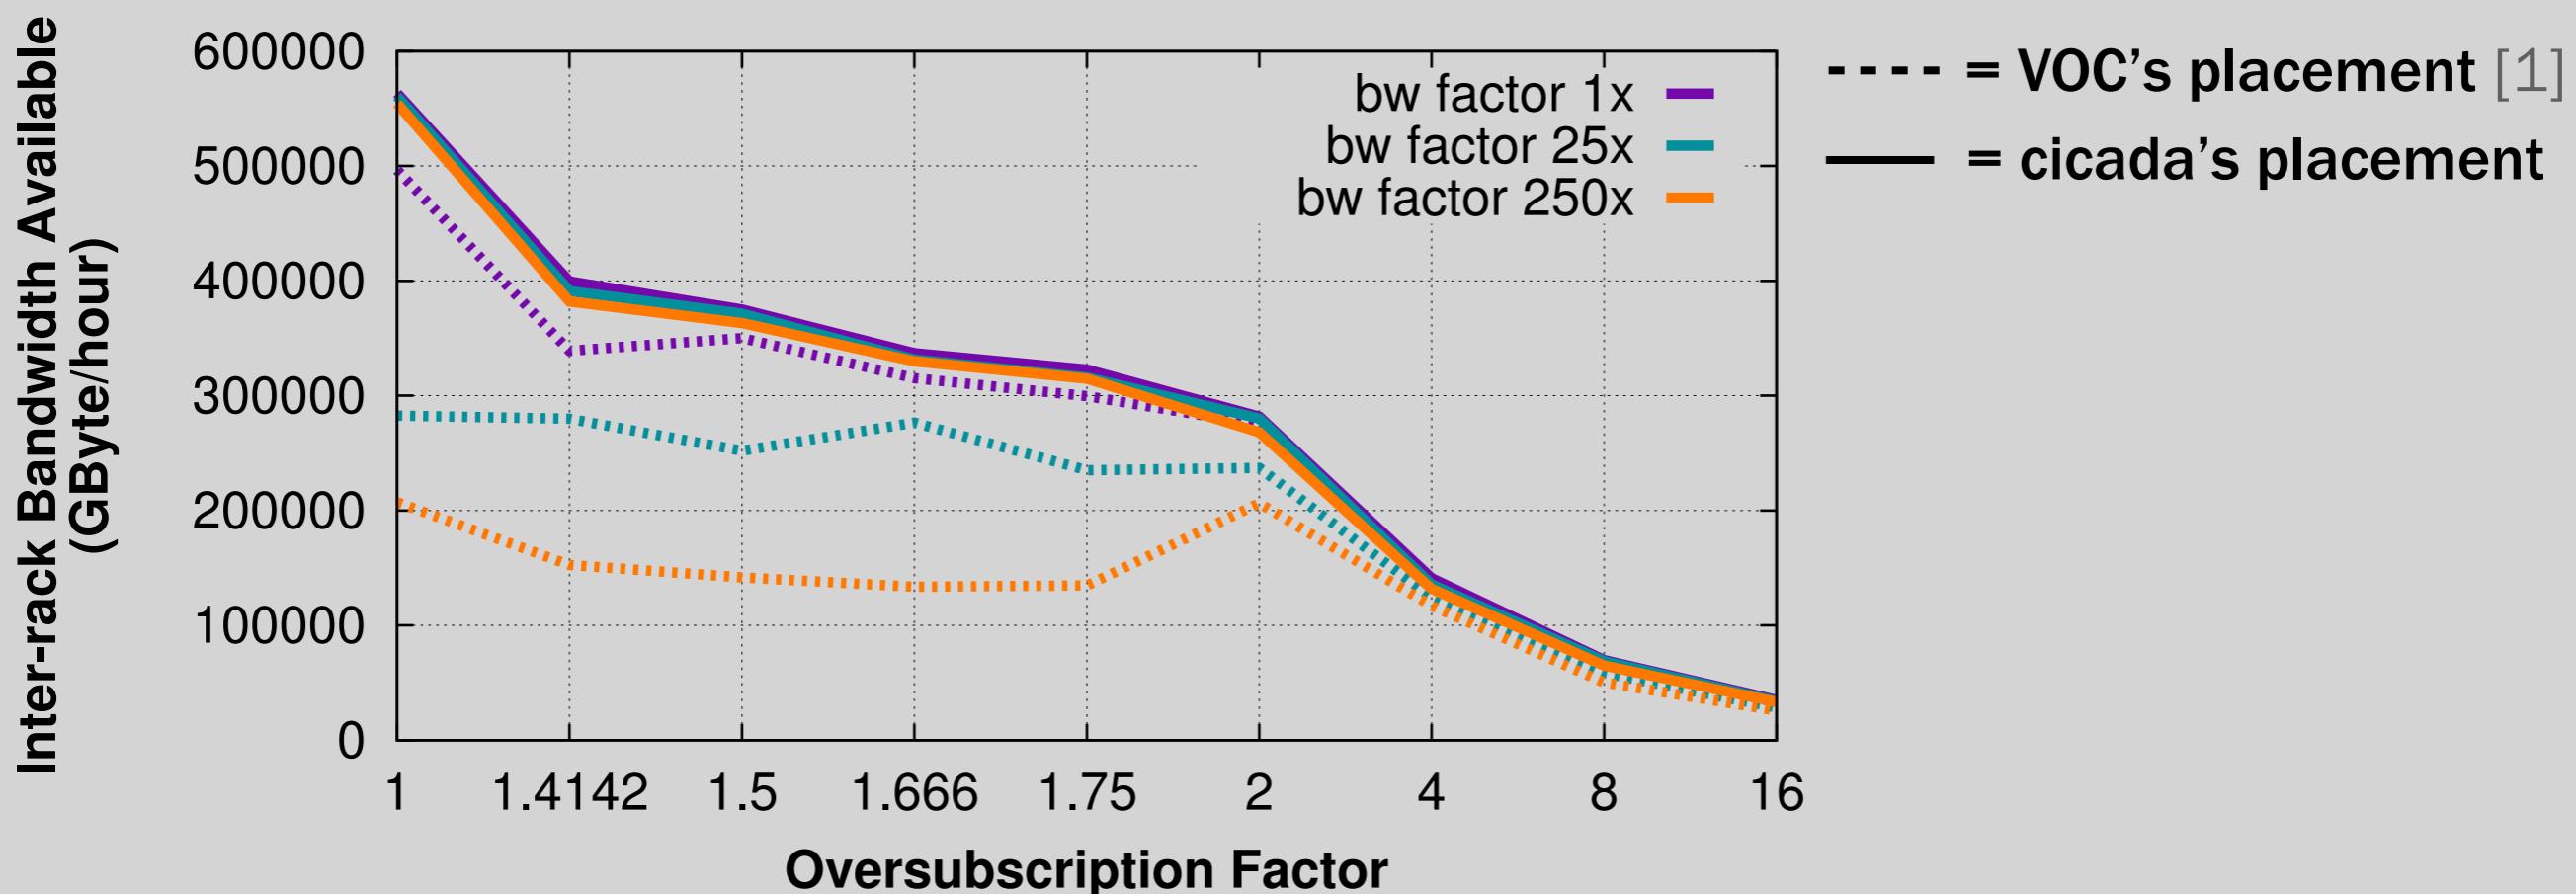
NETWORK UTILIZATION

cicada's **greedy** placement heuristic: place the pairs of VMs that need the highest guarantees on the fastest paths (see paper for details)



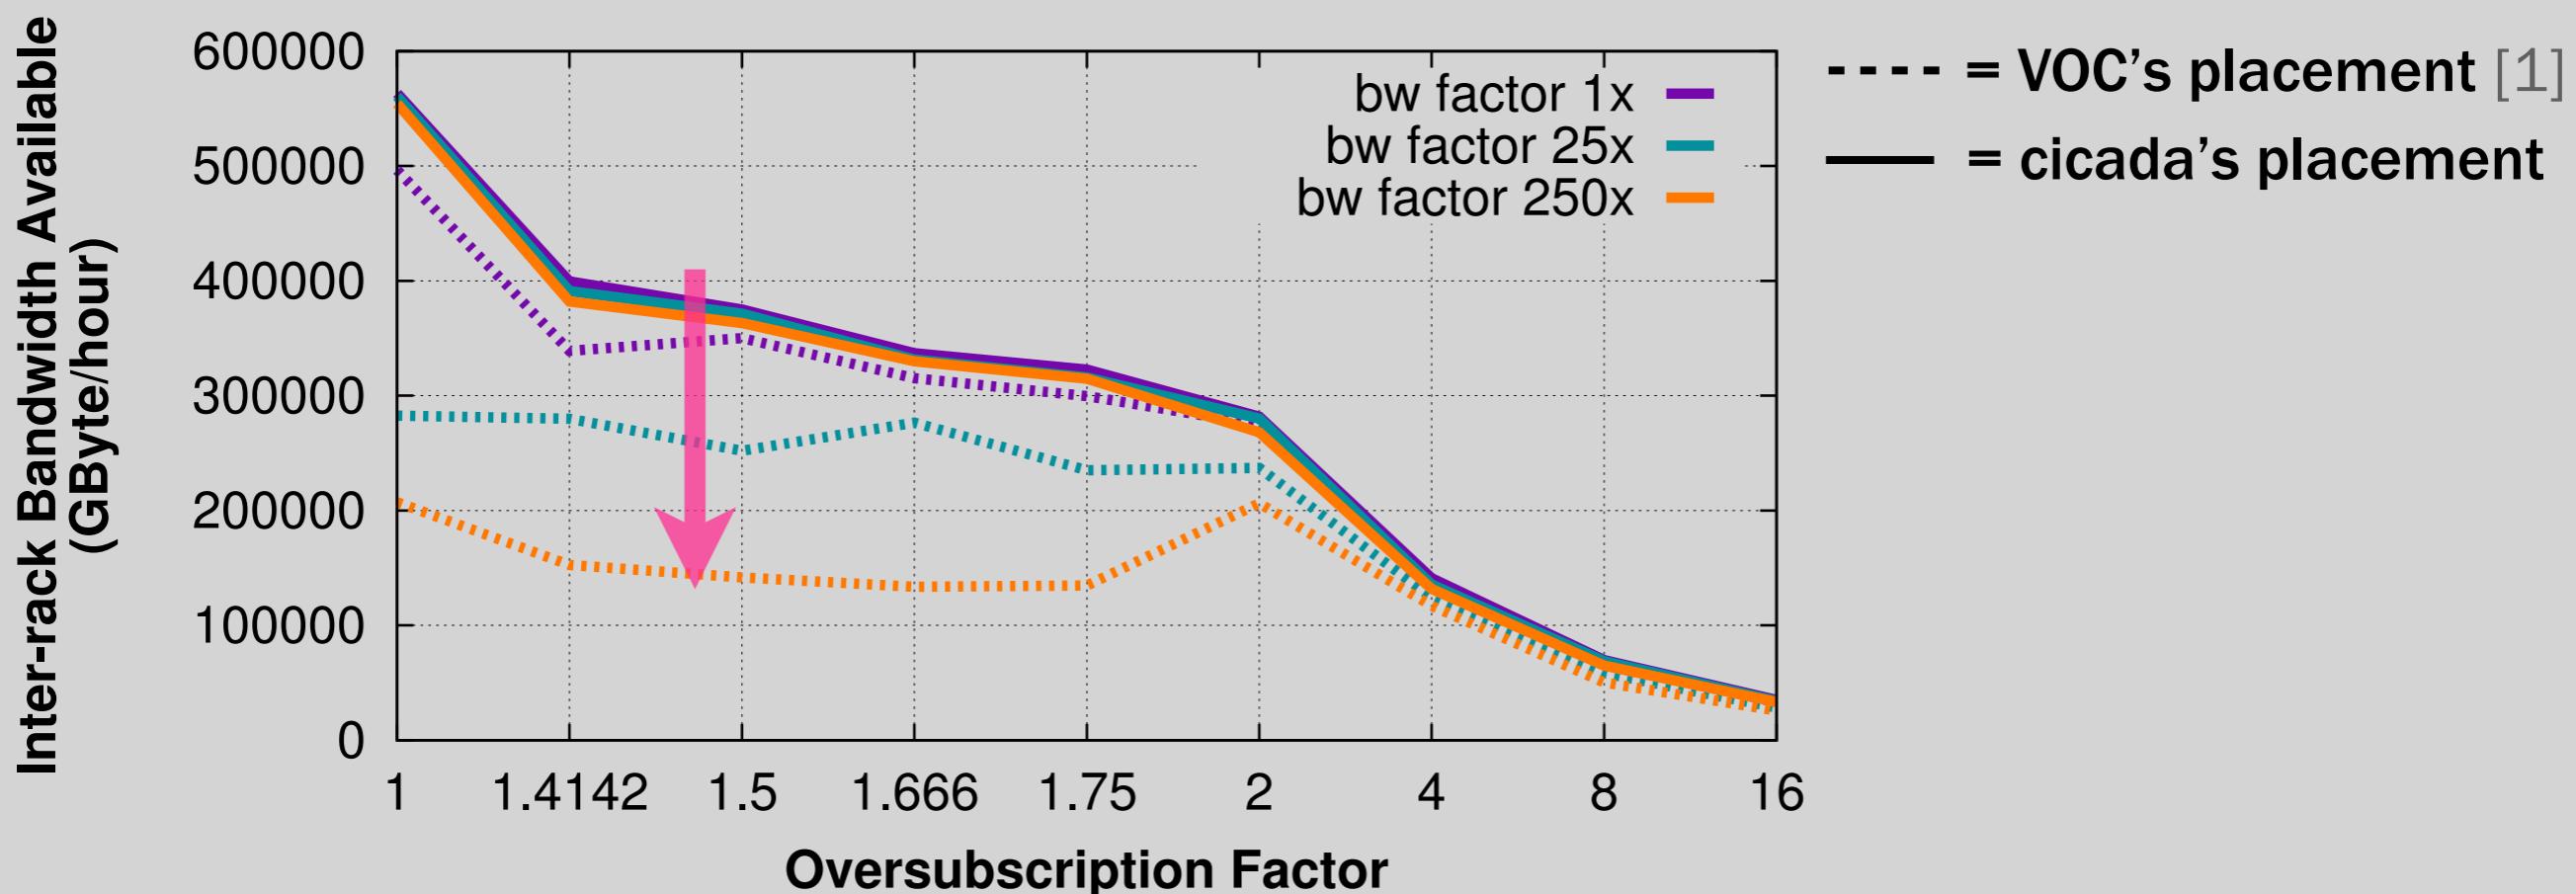
NETWORK UTILIZATION

cicada's **greedy** placement heuristic: place the pairs of VMs that need the highest guarantees on the fastest paths (see paper for details)



NETWORK UTILIZATION

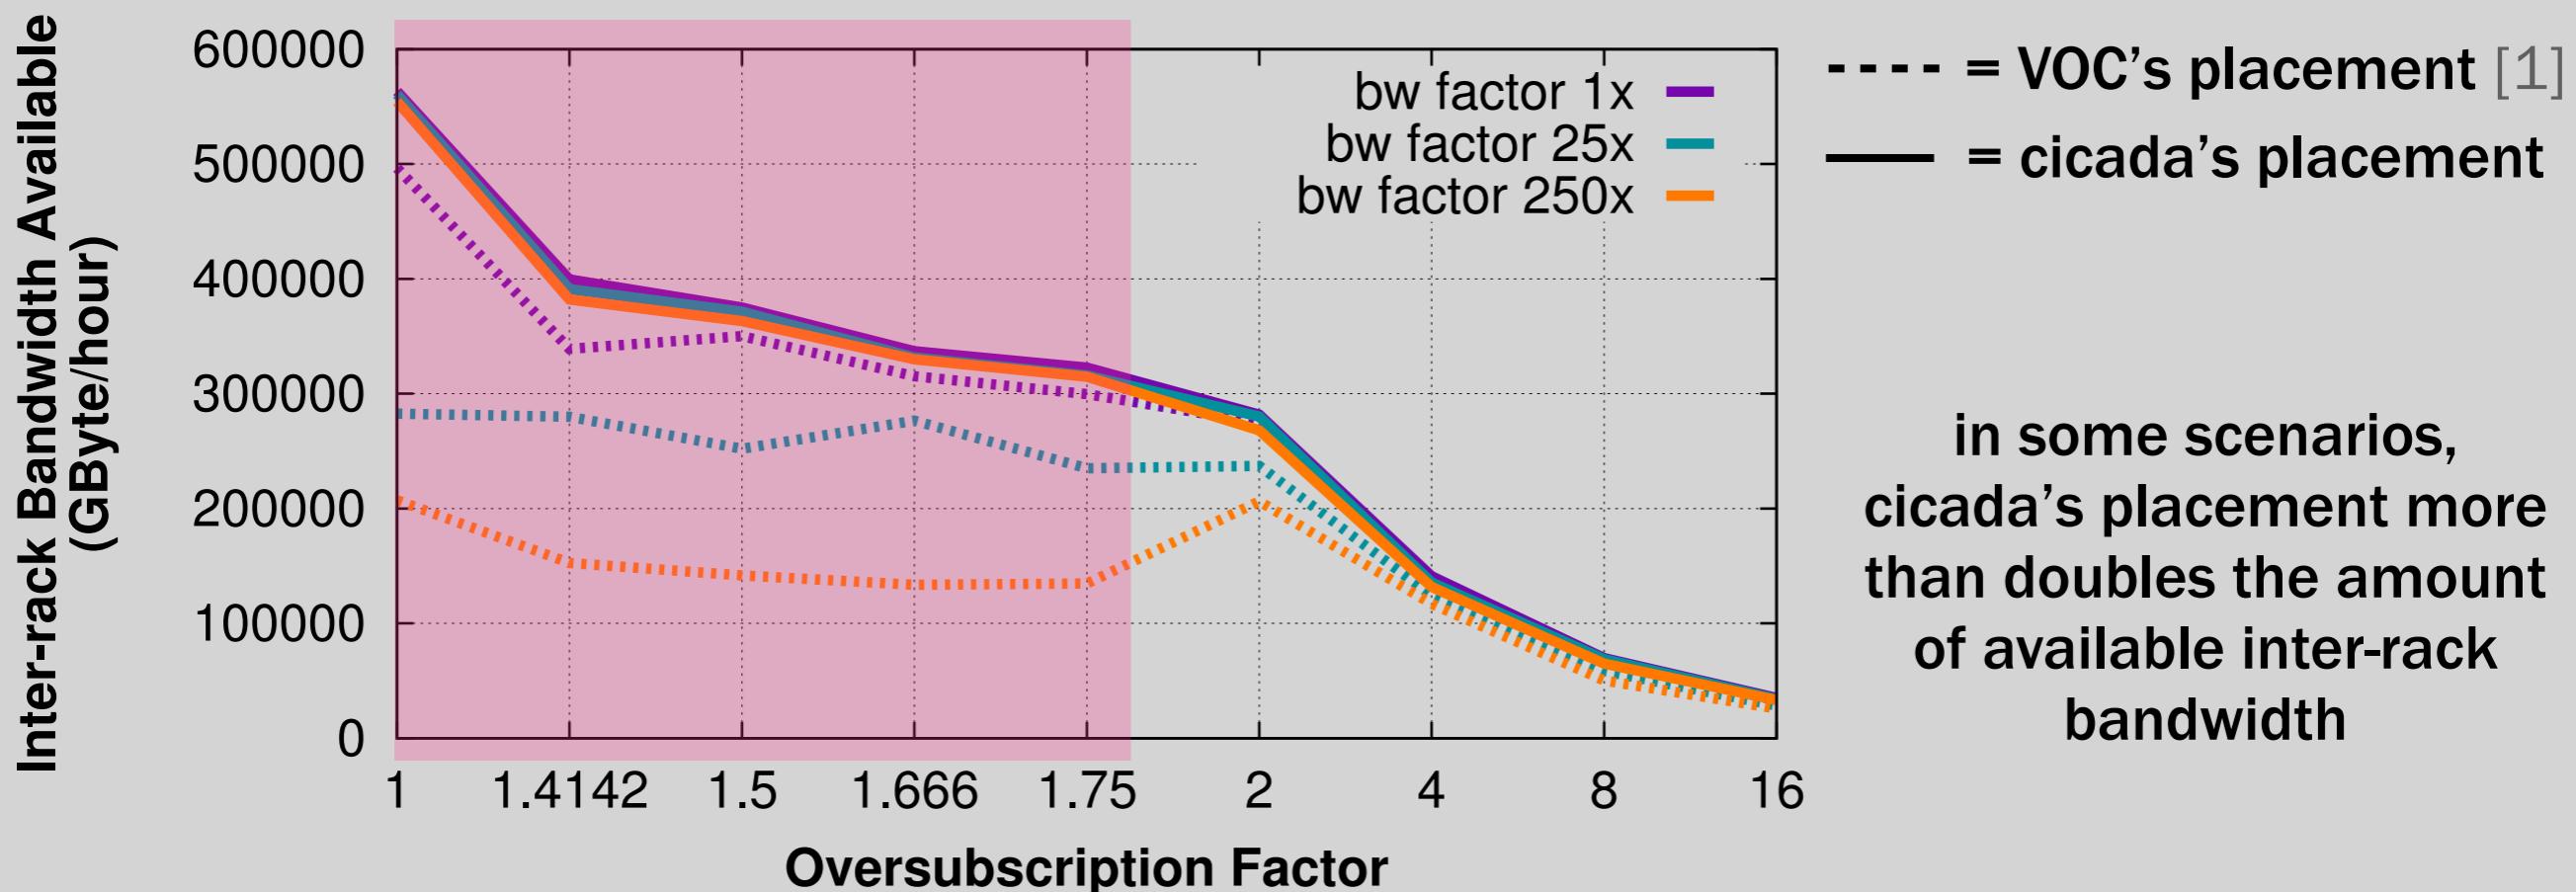
cicada's **greedy** placement heuristic: place the pairs of VMs that need the highest guarantees on the fastest paths (see paper for details)



as applications use more bandwidth, VOC's placement wastes more inter-rack bandwidth

NETWORK UTILIZATION

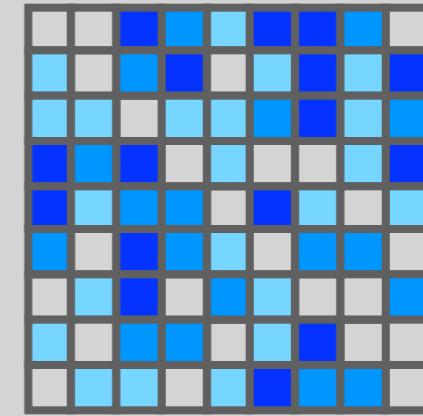
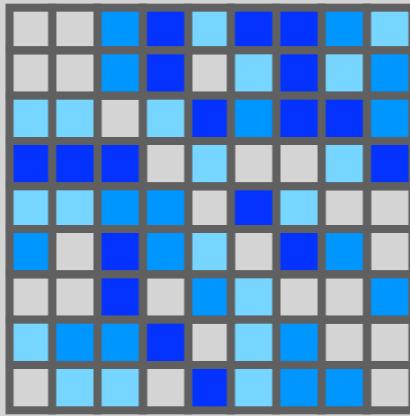
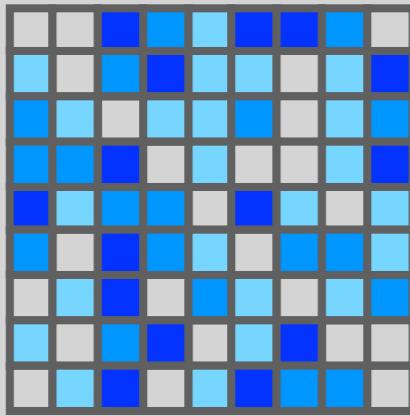
cicada's **greedy** placement heuristic: place the pairs of VMs that need the highest guarantees on the fastest paths (see paper for details)



as applications use more bandwidth, VOC's placement wastes more inter-rack bandwidth

SUMMARY

cloud applications exhibit variability that
existing models don't capture

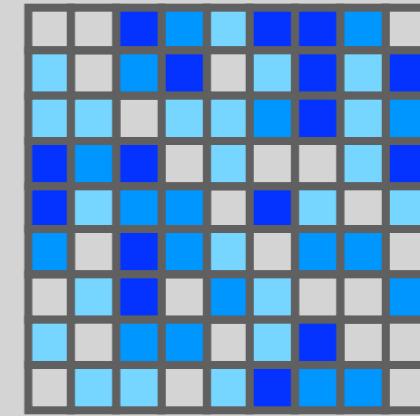
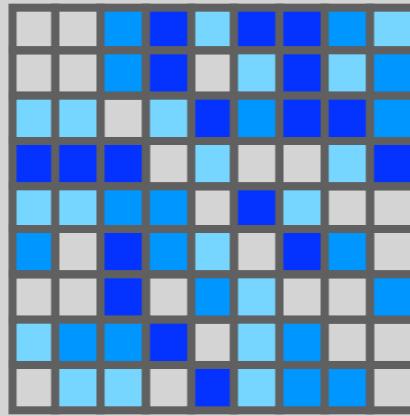
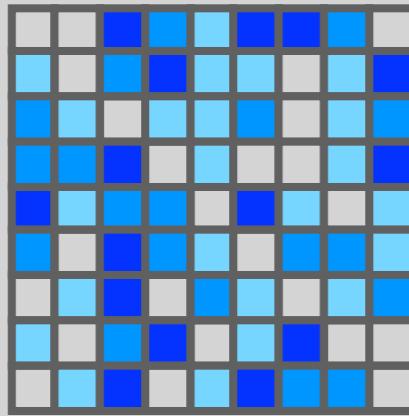


cicada captures this variability, and provides guarantees that are
accurate
calculated quickly*
require little history
increase network utilization

*see paper for this result

SUMMARY

cloud applications exhibit variability that
existing models don't capture



cicada captures this variability, and provides guarantees that are

accurate

improve relative error by 90%

calculated quickly*

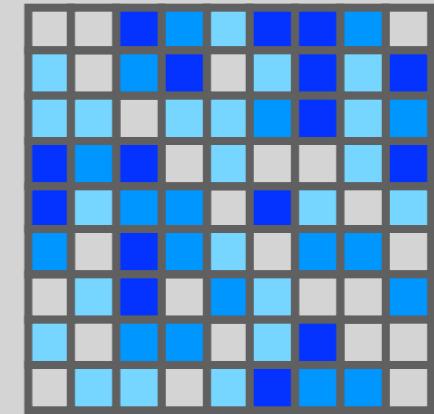
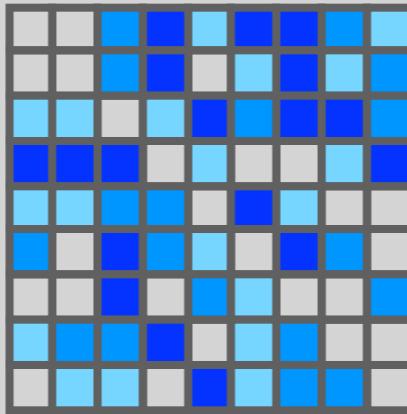
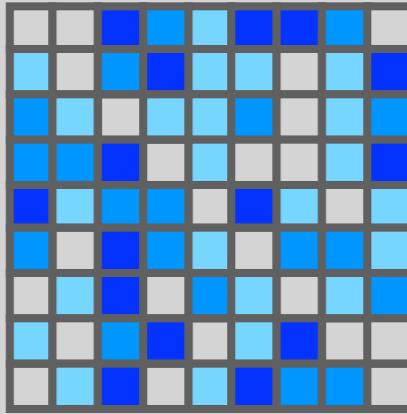
require little history

increase network utilization

*see paper for this result

SUMMARY

cloud applications exhibit variability that
existing models don't capture



cicada captures this variability, and provides guarantees that are

accurate

improve relative error by 90%

calculated quickly*

< 10milliseconds per application

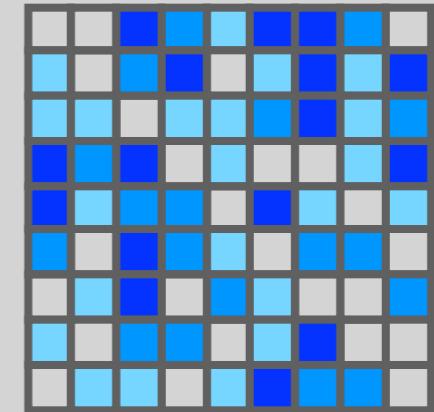
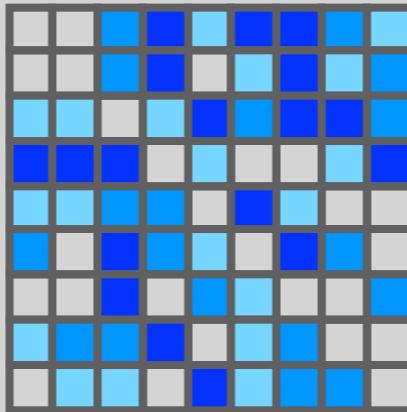
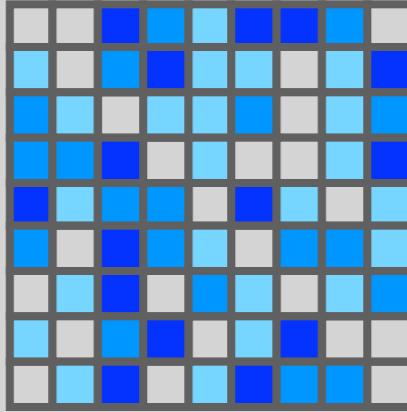
require little history

increase network utilization

*see paper for this result

SUMMARY

cloud applications exhibit variability that
existing models don't capture



cicada captures this variability, and provides guarantees that are

accurate

improve relative error by 90%

calculated quickly*

< 10milliseconds per application

require little history

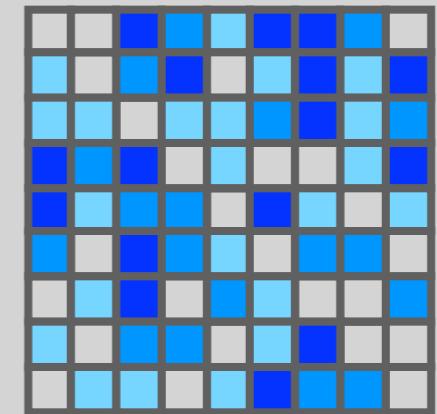
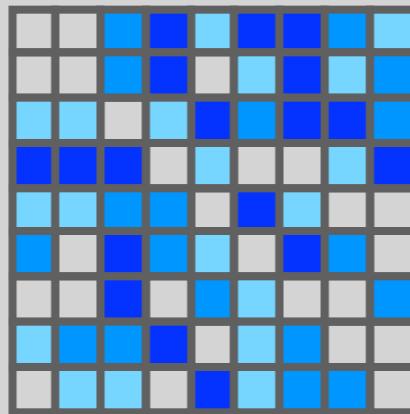
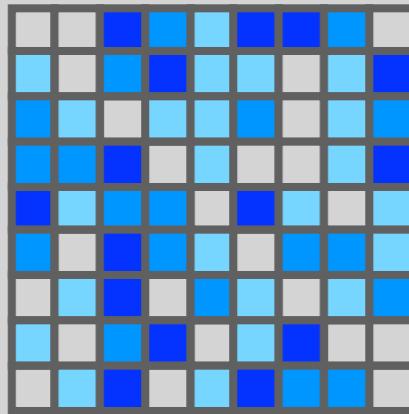
1–2 hours for most applications

increase network utilization

*see paper for this result

SUMMARY

cloud applications exhibit variability that
existing models don't capture



cicada captures this variability, and provides guarantees that are

accurate

improve relative error by 90%

calculated quickly*

< 10milliseconds per application

require little history

1–2 hours for most applications

increase network utilization

in some cases, inter-rack
utilization is doubled

*see paper for this result