

“Tosafot”

Persistent disk storage in EC2: Rashi explained that this means EBS snapshots. Do not get confused and think this means EBS volumes, for EBS volumes themselves are not persistent. The data stored on it only becomes persistent when the volume is snapshotted.

Do not sleep well: Not like Rashi, who says it's because they are worried NFS will break. Rather, they do not sleep well because when the NFS server does break it is very cumbersome to fix because it requires restarting the NFS daemon on all the clients, which causes database downtime.

[Hava Amina – the text's initial position:] Databases running on EC2 instances can be operated identically to the way they are operated outside the cloud. [Asks the Gemara with question #1:] This [seems to] mean that the same procedures and tools that you use outside the cloud will work inside the cloud. But, haven't we learned that the network capabilities of EC2 are limited and do not provide multicast? [Answer attempt #1:] If you don't rely on multicast then you can use Oracle in the cloud. [Rejection of answer #1 with question #2:] Haven't we learned[1][2] that the persistent disk storage in EC2 requires special consideration and care when used for primary database data storage? [Thus, answer attempt #1 is rejected, and question #1 still stands.] Ko Mashma Lan: [we learn from this:] Operating databases inside the cloud requires customized tools. [A further question #3 is posed:] Why shouldn't you consider using NFS? [The answer to question #3 is:] Pok Chazi [go out and see what people are practicing] that the people who use NFS in EC2 do not sleep well. Furthermore, we have a Kabala [a tradition] that HA systems should be spread across AZs, and the additional network latency introduces problems for NFS [especially when used for database storage]. Ko Mashma Lan: [we learn from this:] that existing Oracle solutions do not work well in the cloud. [For further clarification,] What additional “procedures and tools” are referred to [that do not work for Oracle in the cloud]? Application deployment and management frameworks that provide failure recovery for the data tier as well as the application tier, and integrated monitoring.

“Rashi”

Databases: Oracle and MySQL. **Procedures and tools that you use outside the cloud:** More details are provided in the text below. **Multicast:** which is needed by Oracle RAC. **Persistent disk storage in EC2:** EBS snapshots. **Requires special consideration:** And therefore the same tools that are employed outside the cloud cannot simply be used as-is in the cloud. **Customized tools:** that incorporate the special procedures necessary to properly operate the cloud infrastructure. **NFS:** which is used by some databases to share a single storage volume across multiple database instances. **Do not sleep well:** because they are worried that it will break. **HA:** Highly available, fault-tolerant. In the cloud HA systems must provide the HA features in the software layer, without relying on the infrastructure, which has a low SLA. **AZs:** Availability Zones. Systems should be spread across different availability zones in order to avoid system failure when one AZ experiences issues. **Do not work well in the cloud:** because they do not take into account the considerations and limitations mentioned in the text. **Application deployment and management frameworks:** such as enStratus.

[1] <http://ec2ebs-mysql.notlong.com/>

[2] <http://alestic.com/2009/09/ec2-consistent-snapshot>