


## NVMe to NVMe

### Overview

#### Key Features

- ↪ NVM Express compliant
- ↪ 1 to X host interfaces

This architecture proposal is based on NVMe offload IPs. It is integrated in a FPGA, allowing to provide design flexibility (for the PCIe gen interface and the number of host interfaces) and the ability to add processing accelerator IPs for the storage management such as erasure coding and deduplication, encryption...



#### Benefits

- ↪ Ultra low latency
- ↪ Very high throughput
- ↪ Low power architecture
- ↪ Low gate count
- ↪ Reduced time-to-market
- ↪ Easy integration


Such architecture can be configured in many ways depending on the storage management strategy and application requirements. It starts from a basic capacity aggregation: one namespace across the 4 SSDs, until asymmetric : one namespace on one SSD and 10 namespaces on the 3 other SSDs. We can then imagine a lot of combinations allowing the best trade-off in term of redundancy, performance and capacity. It is also possible to imagine multi namespace with different characteristic (encryption, compression...) seen only as one storage SSD. It is also possible to create a simple Raid 1 storage totally transparent for the host software. Based on this simple architecture there are infinite combinations to answer customers's needs.

#### Deliverables

- ↪ Verilog RTL source code
- ↪ Software source code
- ↪ Synthesis scripts
- ↪ Technical documentation
- ↪ Technical support

NVMe to NVMe HBA cards are also an ideal path to computational storage. Integrating computing in the HBA controller (erasure coding, deduplication, encryption, compression) but also advanced computing accelerators such as key-value store, search engine and deep learning

## Description




## Performance

For Samsung NVMe SSD 980 PRO  
 Each NVMe host port is able to reach 3,4GB/s in read. The NVMe device IP can support a Gen3 x16 or Gen4 x8 interfaces, delivering the 12GB/s required to support 4 Gen3 x4 host controllers.

## Related IPs


### NVMe Device IP



The IPM NVMe Device is a verilog IP ready to be integrated in a FPGA/ASIC. IP-Maker has developed a full hardware implementation of the NVMe protocol on the device side. It is based on an Automatic Command Processing Unit and a multi-channel DMA to perform data transfers.

This architecture is easy to integrate with standard interfaces (AXI/Avalon), between the PCIe and memory controllers. This low latency design is ready to support emerging memories, such as MRAM and ReRAM technologies. The IP-Maker NVMe IP is UNH-IOL NVM Express compliant.

### NVMe Host IP




The IPM Host NVMe is a verilog IP ready to be integrated in a FPGA/ASIC. This IP is a full hardware implementation of the NVMe protocol on the host side. It manages the NVMe commands and the data transfer without requiring any software (included system initialization).

This offload engine avoids the use of CPUs, delivering a best-in-class performance/power NVMe host controller. It can be used with any compliant NVMe SSD available on the market.

## Data flow arbitration software

IP-Maker provides a basic data flow management software. It is an excellent starting point to design specific NVMe-to-NVMe HBA.



[www.ip-maker.com](http://www.ip-maker.com)  
[contact@ip-maker.com](mailto:contact@ip-maker.com)  
 +33 972 366 513

Les Néréides - 55 rue Pythagore  
 13290 Aix en Provence / France