

Scaling Up With OpenSIPS

By

Paresh Lukka

Chief Architect and VP of Technology

Crexendo Business Solution

Agenda

Discuss how we built a high performance, scalable and Feature Rich Communication Platform using OpenSIPS

OpenSIPS Performance



Each OpenSIPS Node:

- Quad Core/8GB Memory
- Simultaneous Requests Handled by one node :
 - 5,000 REGISTER + 3,200 NOTIFY + 1,600 ACK + 1,200 INVITE + 600 BYE + 600 CANCEL + 300 INFO + 100 REFER = 12,600 requests/minute



OpenSIPS Performance

- Process Creation: Average 65/sec
- Kernel Context Switching: Average 3,400/sec
- Average Memory Used: 3G
- Average CPU Used: < 3% at peak load

OpenSIPS - Steady Performance



OpenSIPS - A Lot of Room to Grow



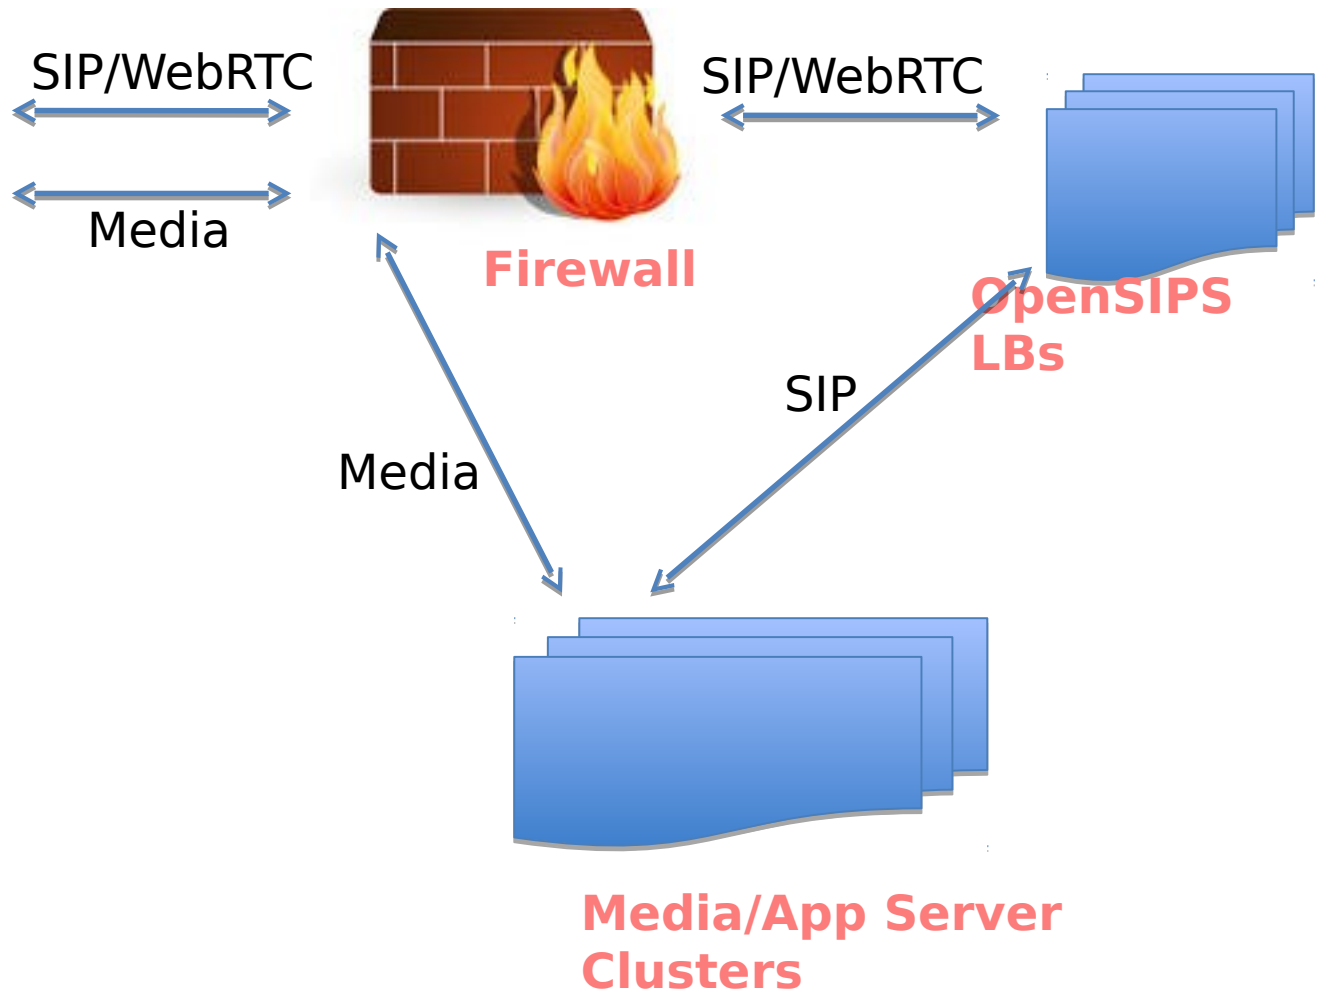
The Minimalistic Approach

- Signaling only – No RTP Proxy
- Registrar
- Topology Hiding
- NAT Helper, NAT Pinging
- TLS Termination
- HA/Failover + Loadbalancing
- Dialog Caching for sticky sessions
- Pike
- Router for Different Application Clusters

Architecture



Endpoints & Carriers



Advantages

- OpenSIPS never a point of failure – Distributed Endpoints, Load Balanced Media/App Servers
- Take advantage of full feature set offered by media/app servers and not involve OpenSIPS with complicated app logic/media handling
- Use OpenSIPS as terminator of various protocol/transport and unify everything behind it – Media/App servers don't need to implement various protocols (like WebRTC) themselves
- Tremendous horizontal scalability – no need for unnecessary POPs
 - Currently one OpenSIPS HA pair can handle 25 media/app servers behind it and a lot more room to grow – not CPU/memory bound at all
- Very easy maintenance/upgrade cycles. OpenSIPS dispatcher mechanism helps drying the traffic to media/app servers. Essentially zero downtime.

Limitations

- I/O
- Log Rotation
- Ping to TLS/TCP endpoint – CPU spikes
- Configurability of Pike per domain absent