SELinux policy for Slurm

Gilles WIBER - CEA
Mathieu BLANC –CEA
M’hamed BOUAZIZ - Atos
Liana BOZGA - Atos

26-09-2017

Bull
atos technologies

SELinux policy for Slurm

- Cyber security in HPC
- SELinux presentation
  - SELinux basics and benefits
  - How it works
  - Challenges
- SELinux Performance Results
- SELinux for Slurm
  - Slurm architecture
  - Confined processes
  - Confined features
- Future work

Cyber Security in HPC

Securing an HPC?

- HPC have become increasingly desirable targets to attackers
- HPC protection includes:
  - Protecting the set of distributed resources (network access, compute nodes, storage…)
  - Ensure infrastructures, users, data, and jobs are running securely
- Standard security must be enhanced to address issues of HPC security

Encountered challenges

- The issues related to HPC security are not exactly like general computer security
  - Addressing and implementing traditional security solutions is the base for HPC (large and heterogeneous environment)
  - Maintaining and monitoring cluster security is a challenge due to large-scale skill requirements and production constraints
  - Keeping performance (or very low impact) is mandatory

Mitigating HPC threats

- Identity and authorization management must be put in place / managed (Kerberos, LDAP, etc)
  - the solution must scale

- Confine and monitor network traffic
  - To maximize computing resources availability

- The HPC should be perceived as one system, not as a set of systems
  - Multi-level security must be put in place (in-deep security)
  - Component security in addition to global security

    => Securing HPC services using SELinux


2
SELinux Presentation

General introduction

- Security-Enhanced Linux released by the NSA
  - integrated into the Linux Security Modules (LSM) framework standard kernel
  - implements MAC (Mandatory Access Control) based security policies
  - provides service and user confinement
- Policy is the heart of SELinux
  - A set of rules determines security and access permissions for everything in the system
  - Defined by Types, Domains, Identities, Roles and Access with associated transitions
  - Expertise is required to write/adapt policies (SELinux, service behavior, system calls, etc)
- Using a policy is simple and doesn't need an expertise in SELinux

Benefits of running SELinux

- Reduce vulnerability against privilege escalation attacks
- Can be used to enforce data confidentiality and integrity control
- Provide fine-grained access control

- To reach that goal all processes and files are labeled with a specific type

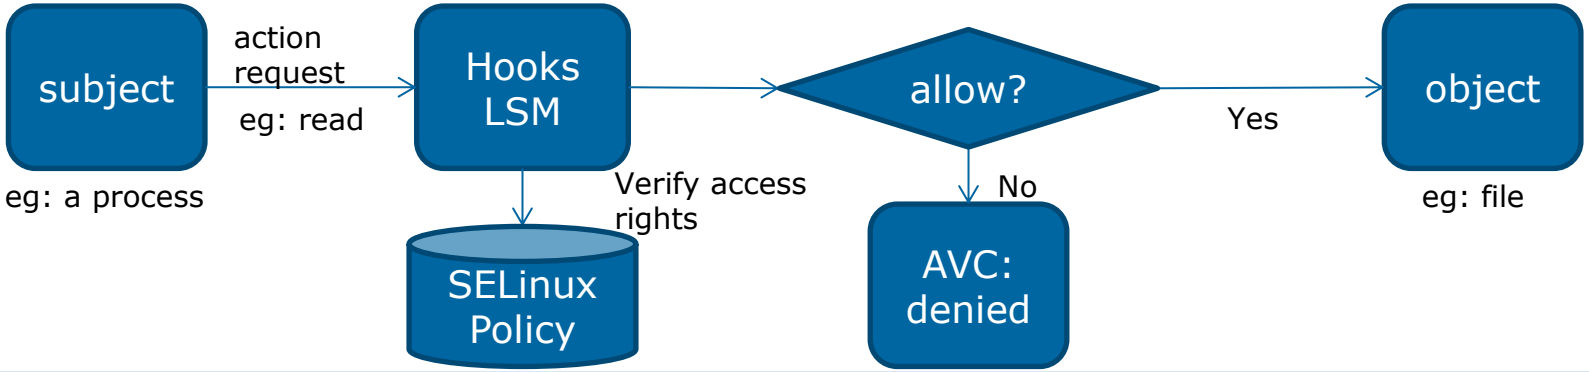

Using SELinux in a HPC?

- Red Hat 7.x provides SELinux targeted policies for standard UNIX services

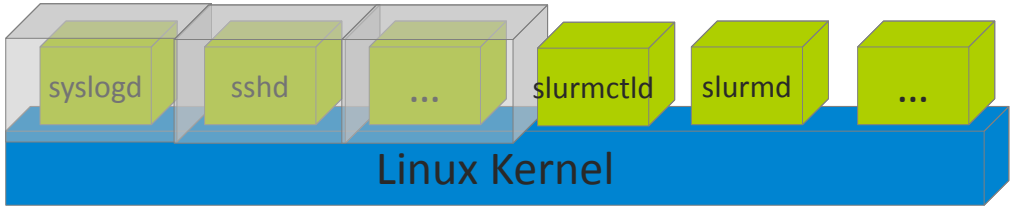

- In a HPC context, resources are distributed
  - SELinux protects local resources for each node
  - A global policy will be loaded even if all services are not installed


3
SELinux Performance Results

Suite tests description

Name	Type	Parallelism	IO
MPI_Init	MPI	MPI	Aucun
MPI_STREAM	Memory	MPI/OpenMP	Aucun
IMB	MPI	MPI	Aucun
HPCC	MPI	MPI	Aucun
OpenMPBench	OpenMP	OpenMP	Aucun
SHOC	GPU	CUDA/OpenCL/MPI	Aucun
IOR	I/O bandwidth	MPI	POSIX/MPIIO/HDF5
b_eff_io	I/O bandwidth	MPI	MPI-IO
metarates	I/O bandwidth	MPI	POSIX
mdtest	I/O bandwidth	MPI	POSIX
Abinit_Compilation	I/O bandwidth	Aucun	POSIX
SPECViewPerf	Visualization	Aucun	Aucun
Abinit	Compute code	MPI	POSIX/MPIIO
WRF	Compute code	MPI	NETCDF/POSIX
NAMD	Compute code	MPI	POSIX

Compute codes

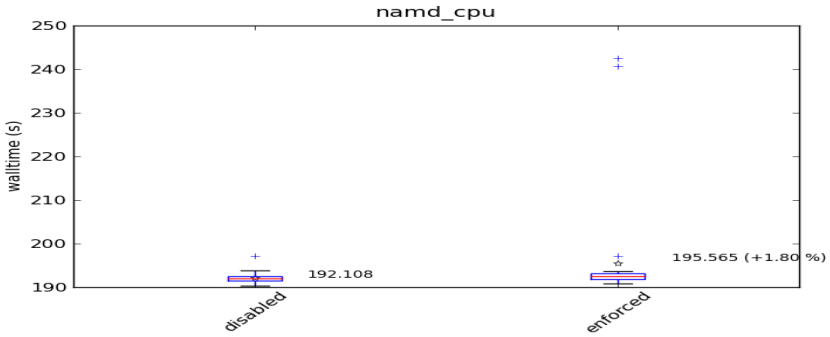

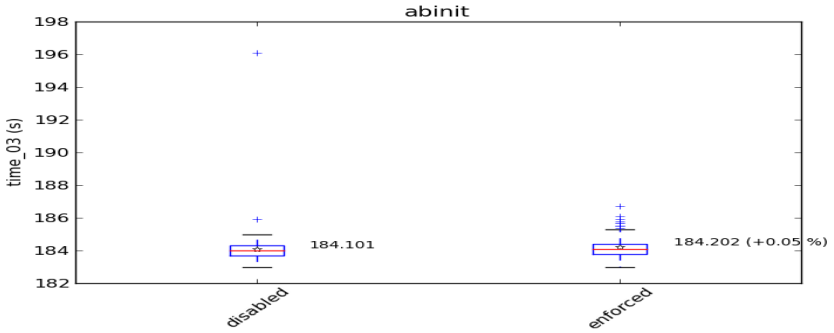

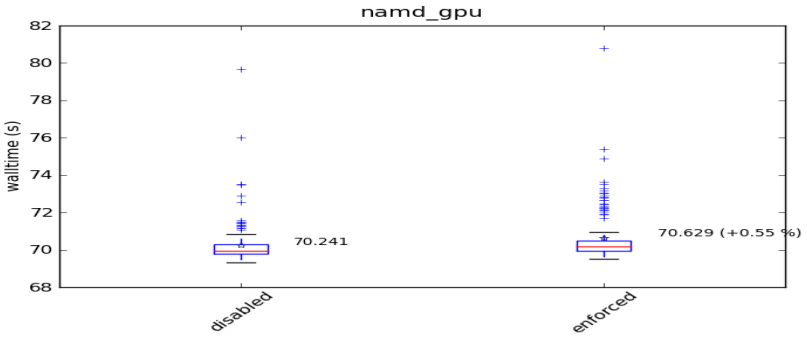

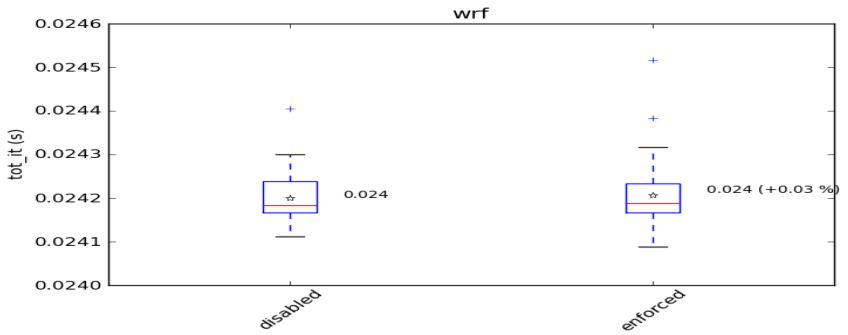

Lustre IOR tests (write bandwidth)

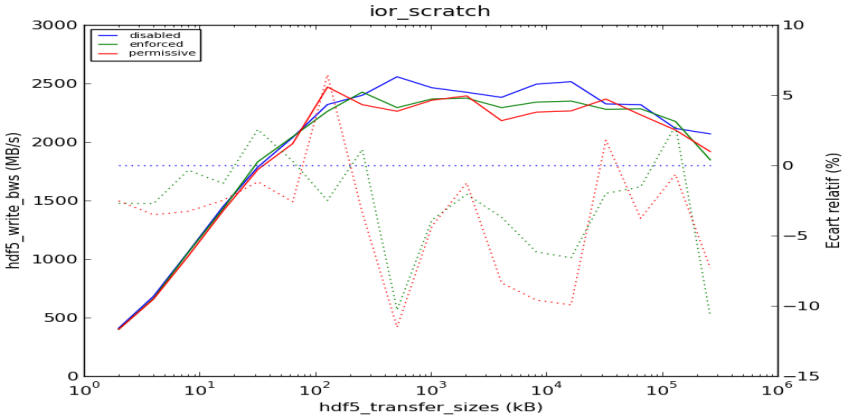

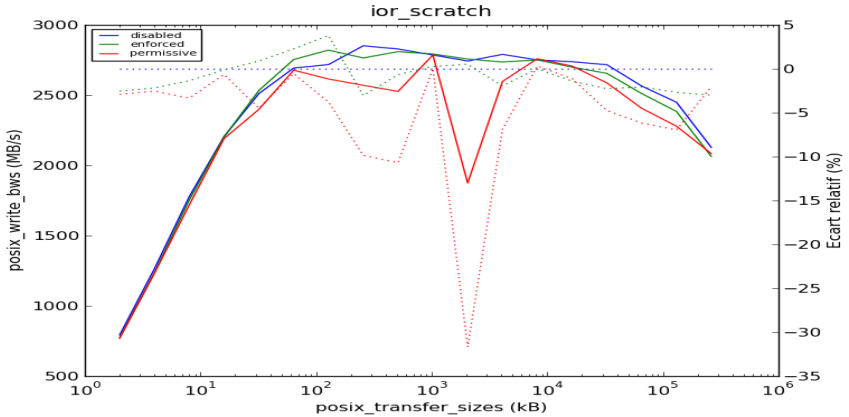

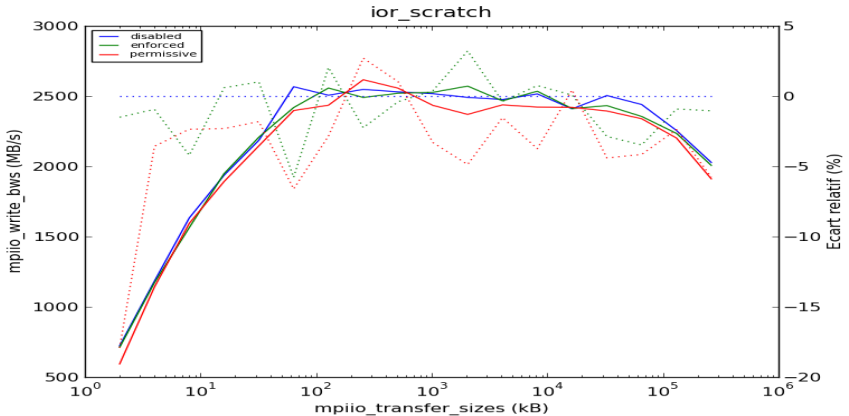

mdtest: files creation

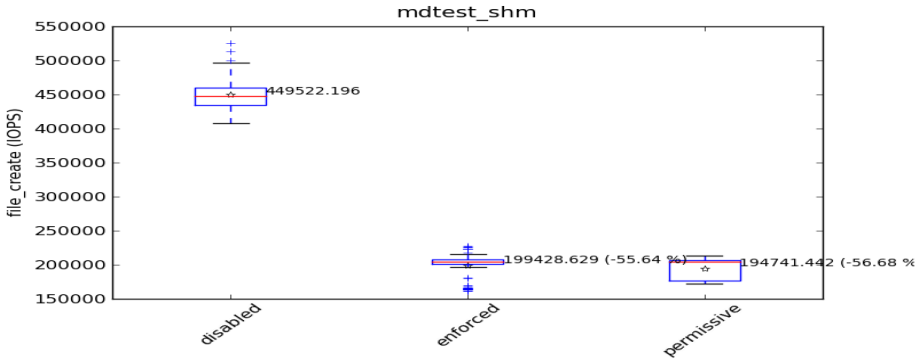

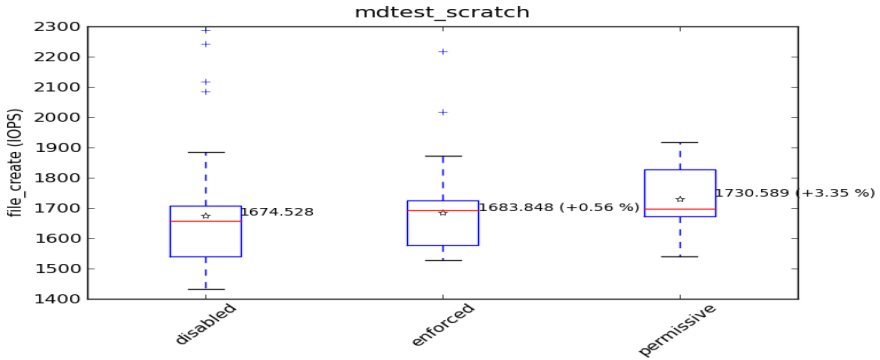

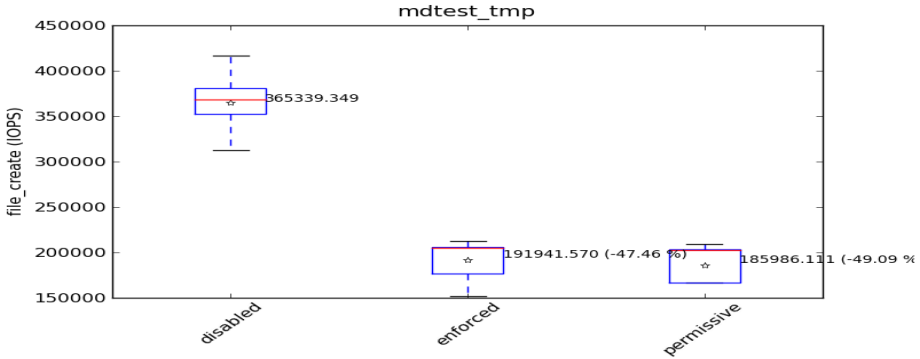

mdtest: files stat

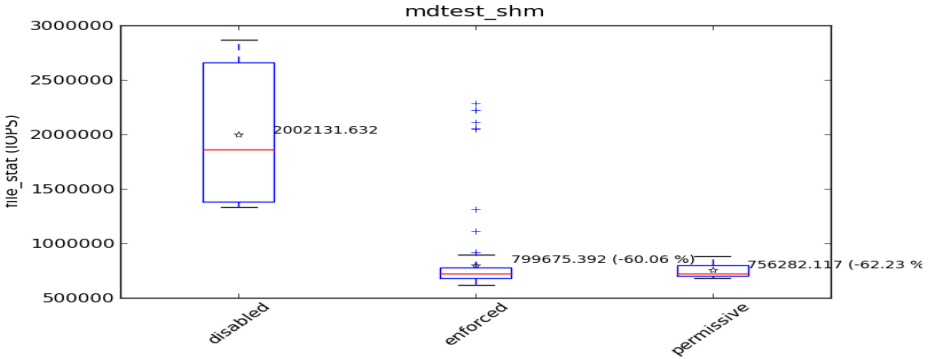

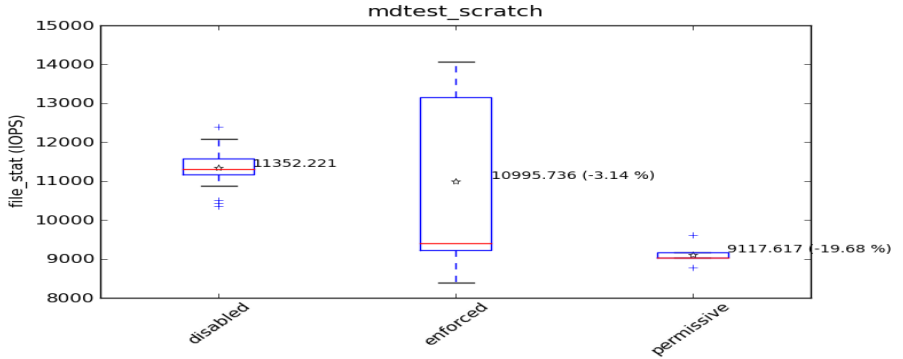

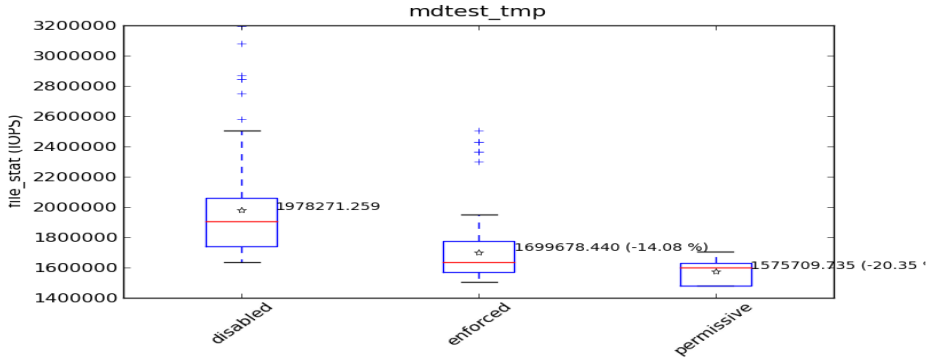

HPCC: latency and bandwidth

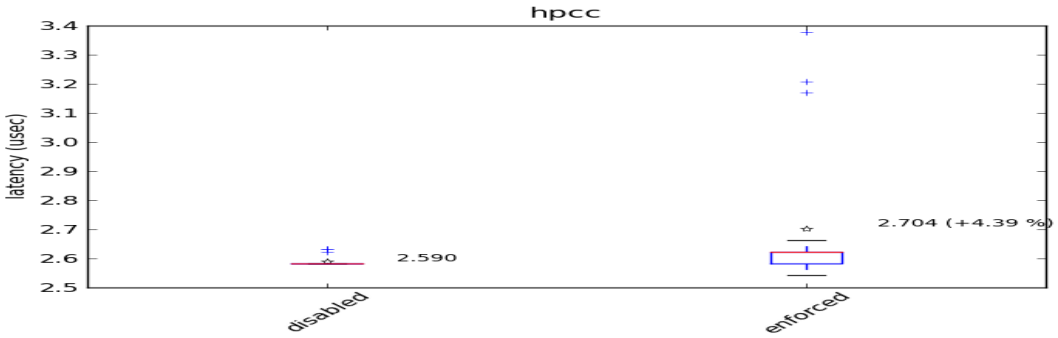

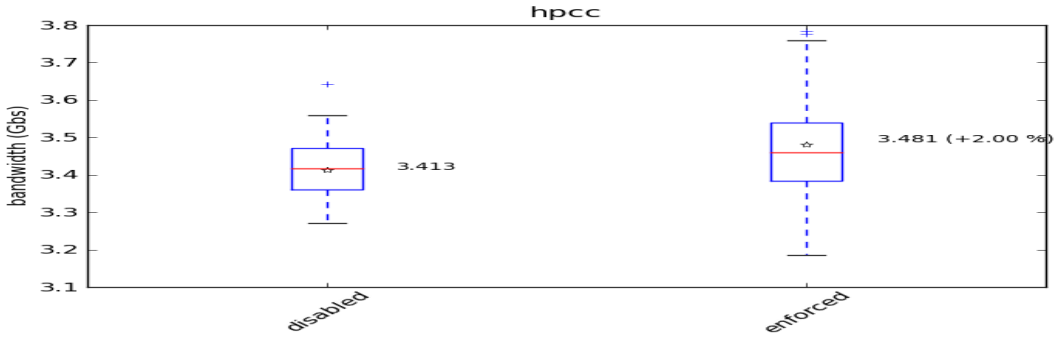

Performance and results

- SELinux impact:
  - No impact on pure compute code (even with GPU)
  - Average 3% to 4% degradation on latency
  - I/O:
    - No impact on I/O bandwidth
    - More than 100% degradation on metadata management

Notice:

Perhaps a good way to limit metadata access !!!

SELinux for Slurm

Security vision on Slurm

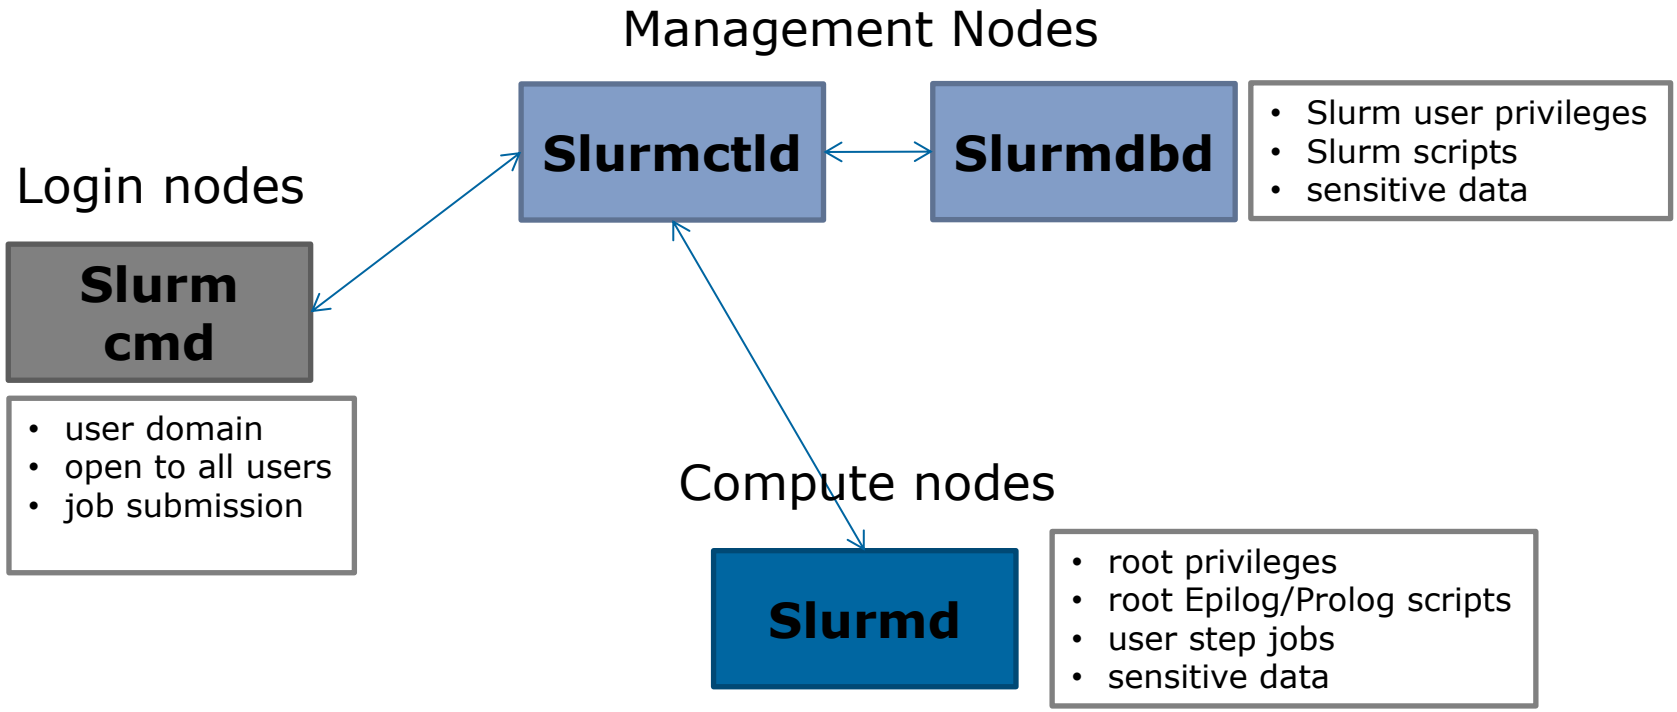

Slurm Policiy definition

- Securing Slurm Using SELinux requires:
  - Confining Slurmd
  - Confining Slurmctld
  - Confining user commands
  - Confining Slurmdbd
  - Confining Slurm scripts

- Confining services => control accesses to local resources such as:
  - network ports, files, directories…

- Writing the policy mustn’t affect the work of Slurm -> The policy must ensure that all features of Slurm are preserved

Confining Slurmd
A view on used resources

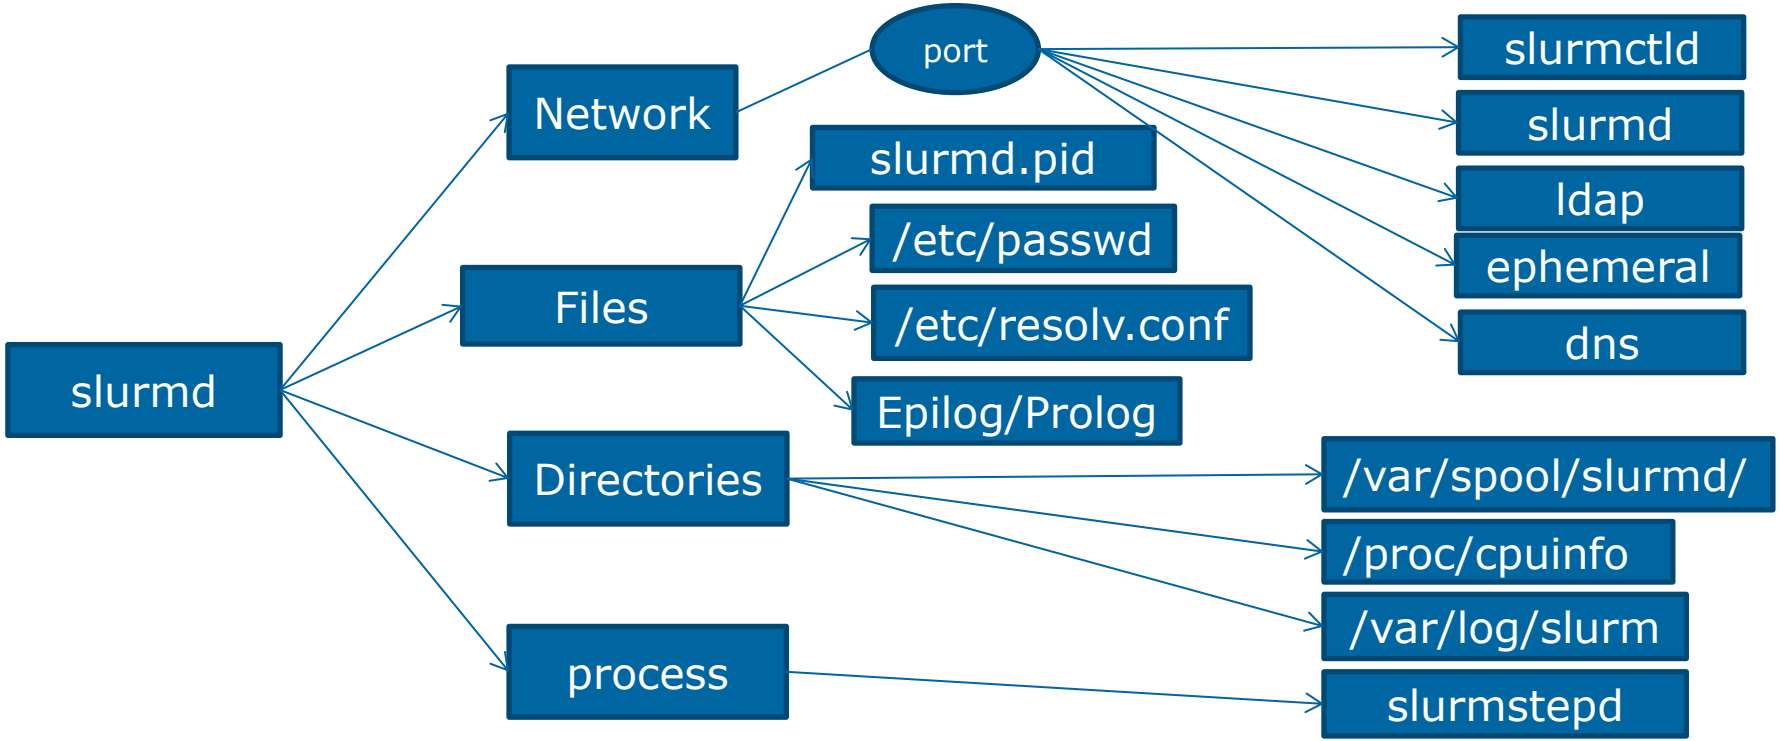

Slurmd Domain
Creating Slurmd Policy

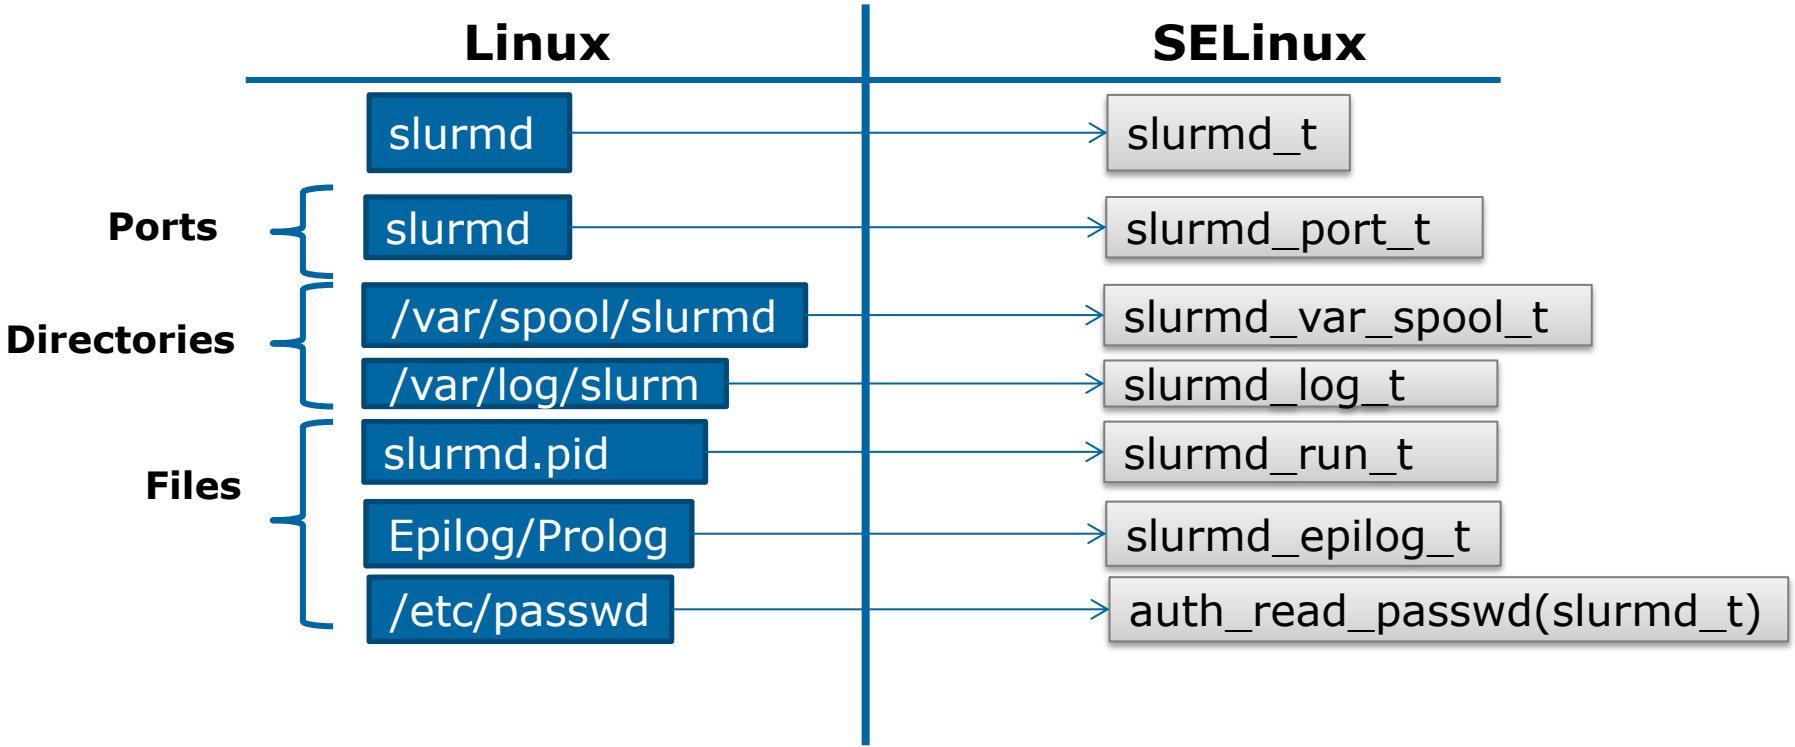

Transition from slurmd_t to slurmstepd_t

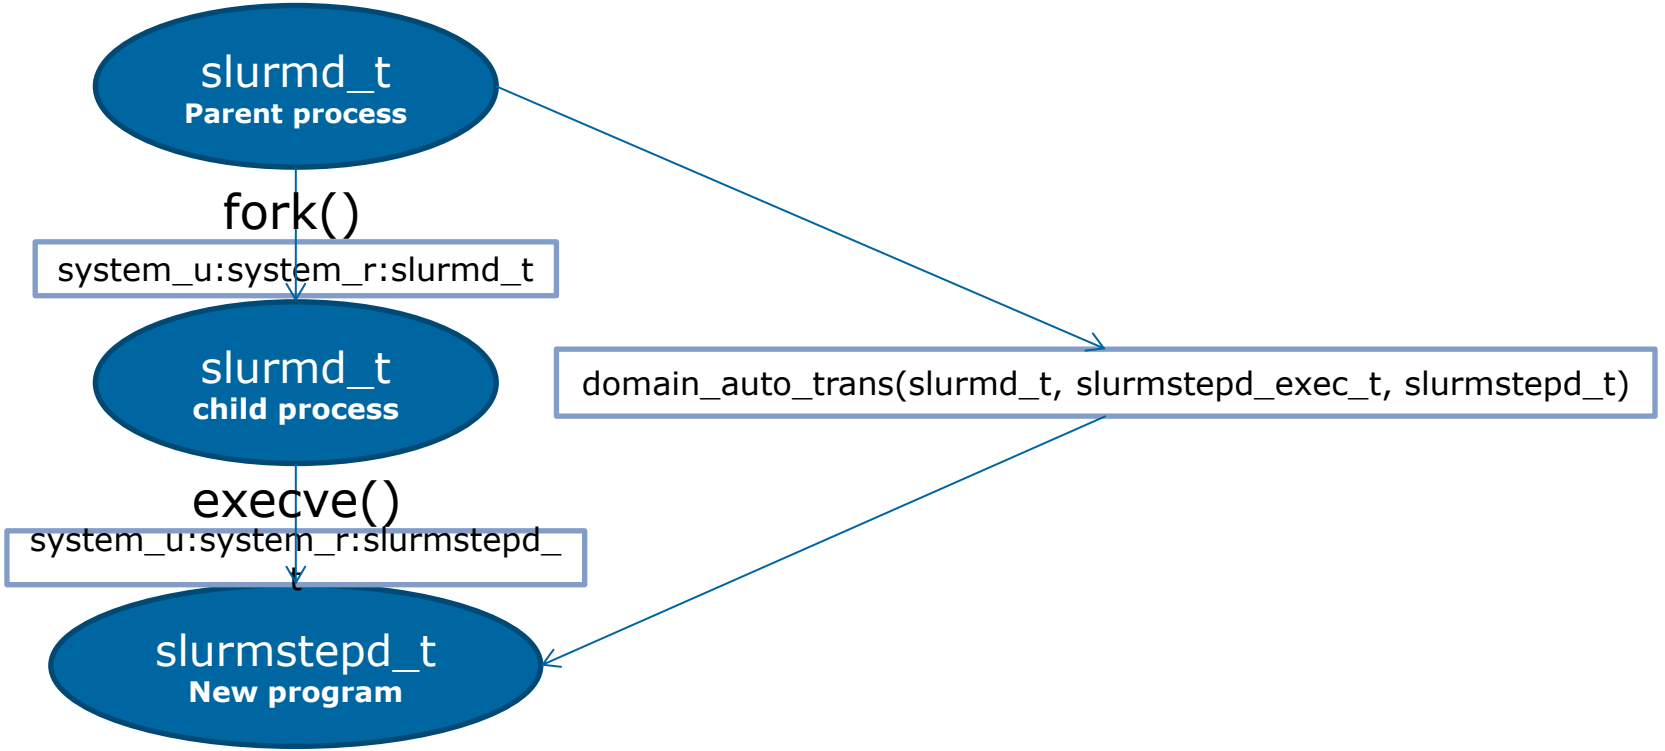

Epilog and Prolog scripts

- Prolog/Epilog requires various privileges depending on job requirement
  => to ease implementation, transitions has been implemented (epilog_t)
  => ”open” environment to execute specific actions outside Slurm policy

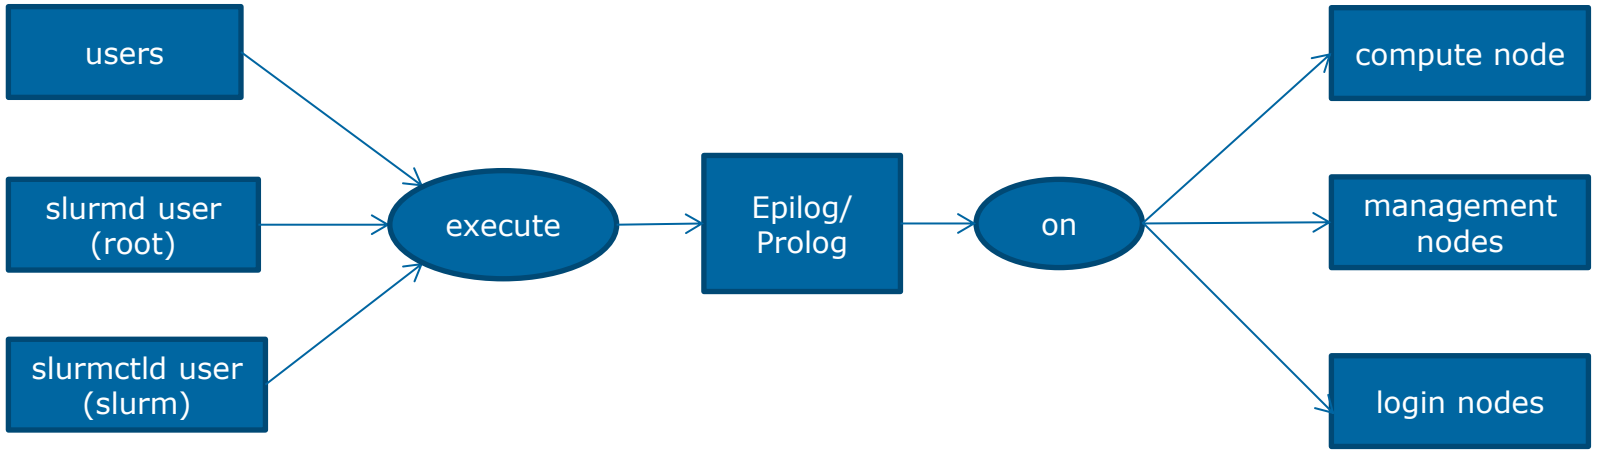

Slurm controller Domain
A view on used resources

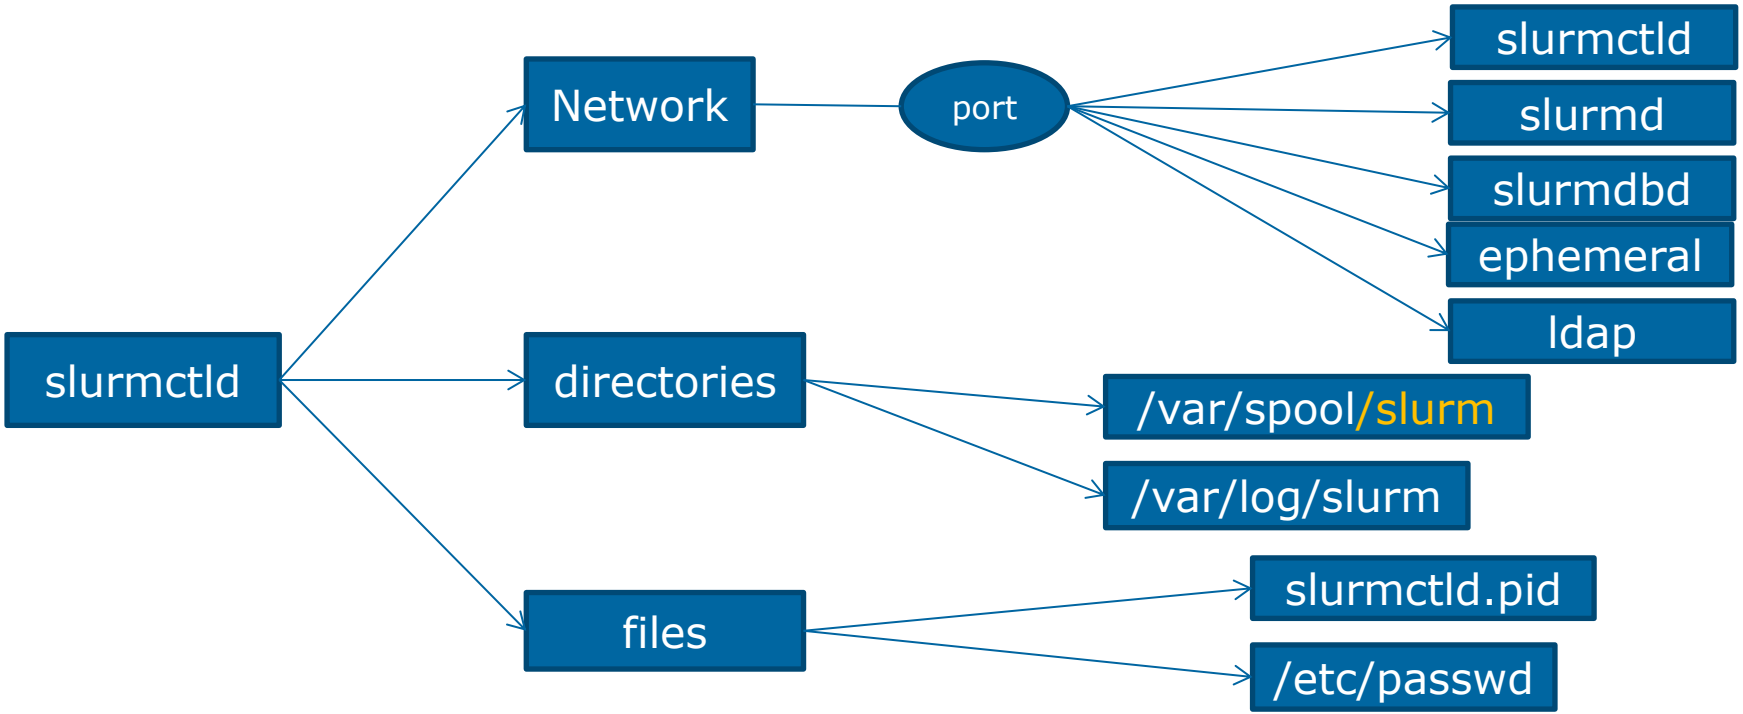

Slurm controller Domain

Creating SELinux Policies

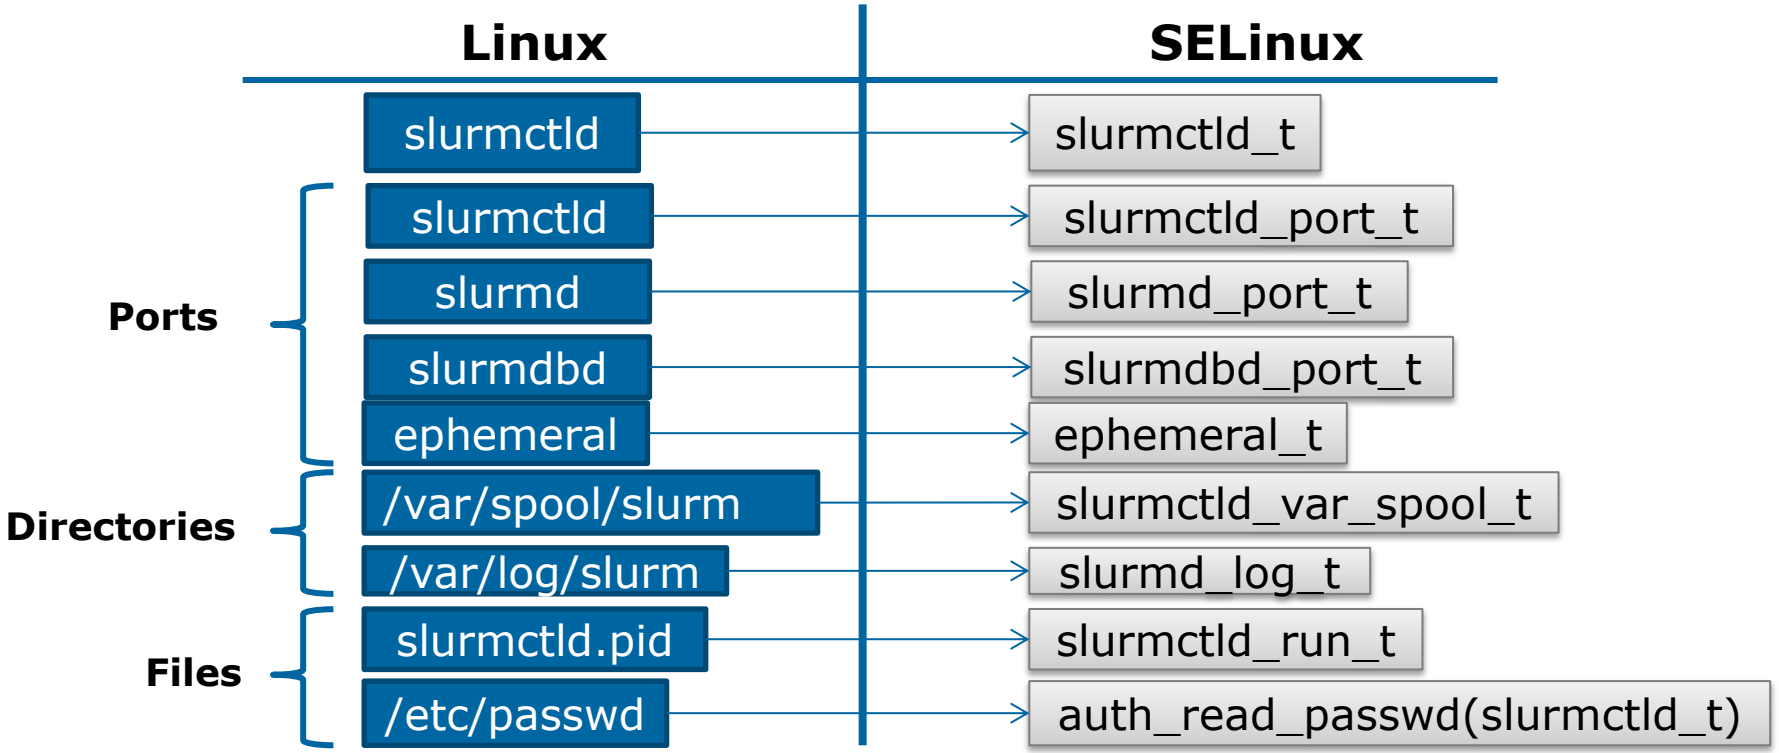

Confining slurmdbd

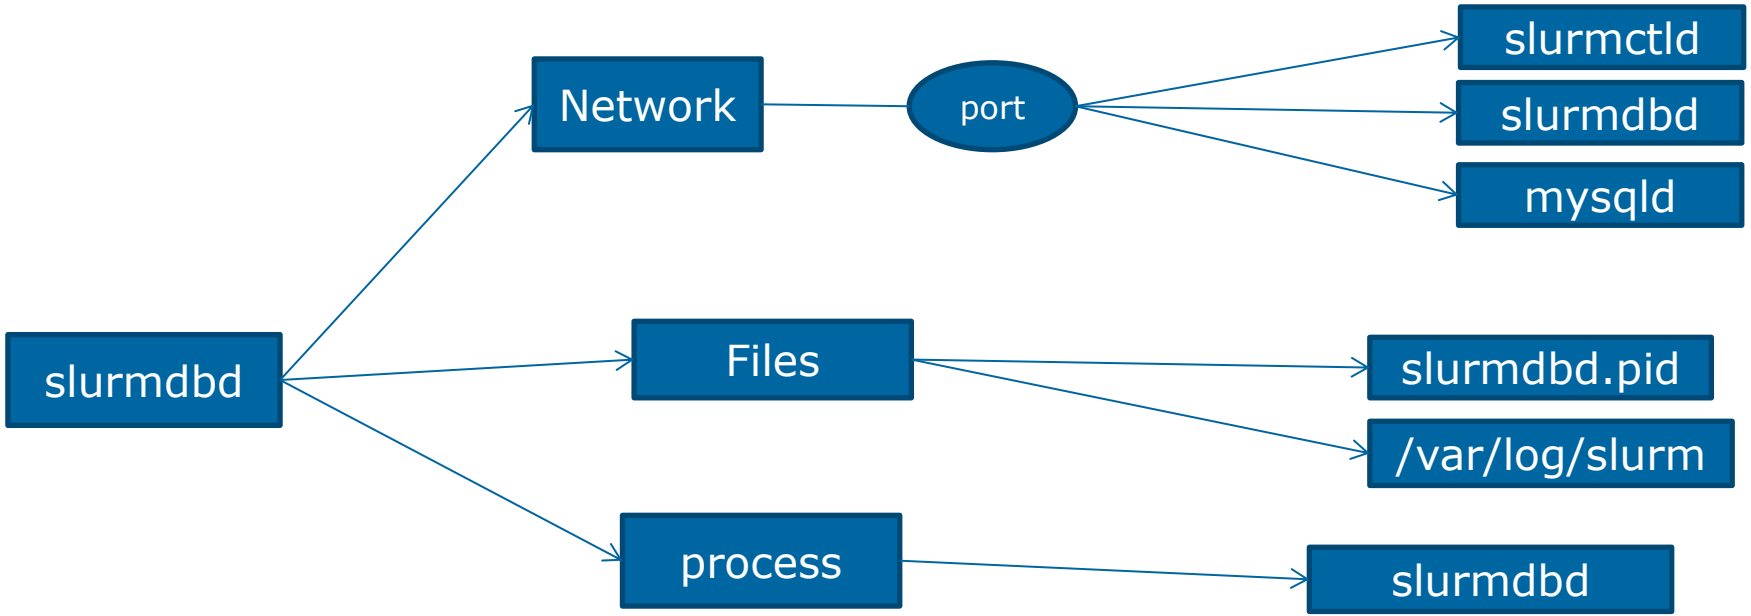

Confining Slurm user commands

- We defined a policy to confine:
  - srun, sinfo, sacct, sbatch, scancel…
- Each command runs in a slurm_t domain
- Malicious users can’t use compiled commands
  - copied or hacked commandes without label
- It allows user cmd to access only authorized Slurm ports

Conclusion

To conclude…

- SELinux Slurm policy can be used to enforce security without additional complexity (pre-defined for Red Hat Linux)
  - Policy also supports some features and plugins such as: X11 spank plugin, interactive jobs, etc
  - Additional work has to be done to extend coverage
- SELinux security provides strong protection against Slurm processes threats (privilege escalation, etc) and also on data integrity (database, accounting, etc) without any additional impact on performance
- But keep in mind that:
  - It is not an all-in-one security solution (part of the a global security design)
  - Policy development and update requires tough expertise

Thanks

For more information please contact:
m-hamed.bouaziz@atos.net