PowerEdge Product Group

Direct from Development

PowerEdge MX Secure Chassis Management

Tech Note by:
Josh Pennell
Michael Brown

SUMMARY

This paper is a discussion of the innovative new security features that are built-into the new Dell EMC MX7000 chassis.

We will cover the secure boot features built into the Management Modules and iDRAC, the ground-up security design incorporating SELinux and least-privilege processes, and our new mechanisms that ensure the security of all management traffic inside the chassis by authenticating and authorizing every component in the chassis, as well as the encryption for all internal management network traffic.

The intent of all of this work is to make a more secure system for customers, one that customers can trust and rely on, and is secured using the best available security techniques against hacking.

Secure Boot

The first principle of security of an embedded management controller is answering the question of what code is running on that management controller. Is the code running on the management controller authentic code from Dell, or has the device been attacked or compromised in any way? The way that we have comprehensively addressed this question for the MX7000 Management Module is our secure boot and chain of trust. Using these techniques, explained below, we can ensure that the Management Module is running unmodified code that has been authenticated by Dell, and that there is no way an attacker has tampered with or replaced any code, either through a supply-chain attack, or through any kind of online attack.

The technique that we use to secure the Management Module is based on the “Chain of Trust” concept. In this concept, each stage of the boot process uses digital signatures to cryptographically verify that the next stage of boot is signed properly before jumping to the next stage.

The beginning of this chain of trust starts in the factory, when the iDRACs and Management Modules that make up the MX7000 are being built. Our hardware is programed with keys, fused to the device, that allow the processor to verify the bootloader prior to starting, the bootloader in turn has its own keys to verify the verify the kernel. Once booted the Kernel runs on a read only file system, further preventing tampering. Each Management Module and iDRAC is also programmed with device unique Identity certificates, which are a public/private keypair used by the device to identify itself as authentic Dell to others. These are signed by the Dell Certificate Authority, are unique to the device and stored encrypted by the devices Hardware Root Key, described below. So each layer of the system verifies that the next is authentic and has not been modified, creating a complete chain of trust from hardware to running code.

One critical part of the secure boot process is the presence of a unique-per-machine Hardware Root Key (HRK). This symmetric encryption key is physically fused into the microprocessor during manufacturing. This HRK is never visible or extractable from the OS or applications running on the management controller, however, applications can make cryptographic requests to the hardware crypto accelerator to encrypt or decrypt information using this key. More importantly to our security design, access to this HRK can be disabled at runtime in a manner that cannot be re-enabled without a power cycle. If at any point in the boot process the system detects that it is running non-Dell code, the HRK is disabled. Why this is important will be explained a little bit later.

SELinux and Least Privilege

After you have ensured that the physical flash is secure and running clean, signed Dell EMC firmware images, the next step is to protect the system from online attacks that would allow an attacker to gain access to a running system through some software vulnerability. We have done this through a combination of two techniques that we have integrated into our development process. First, we have adopted SELinux for all of the management controllers in the MX7000: the Management Module as well as the IDRACs. The 1.0 version of MX7000 firmware ships with SELinux fully enabled and set to the highest level of enforcement out of the box, without any configuration requirements for customers to worry about. The second major runtime security initiative that we have delivered is “least privilege”. This security concept enforces each process to run with an individually unique, non-administrative Unix user account. The combination of these two security techniques helps to mitigate any vulnerabilities that might be found: what would formerly have been a major security hole can sometimes be mitigated down to a minor inconvenience.

This combination of SELinux and “least privilege” protects the sensitive areas of the MX7000 iDRAC and Management Modules. Only the processes associated with establishing machine to machine trust have access to the private key information on the device and access to these processes limited. With defined separation of tasks and access between the difference processes, an attacker of the MX7000 would find their ability to modify or control a system extremely limited, and accessing a remote system unlikely.

Machine to Machine Trust

So far, we’ve been building up a single system piece-by-piece in a secure manner, first by encrypting and verifying the boot process, next by ensuring that the running firmware image is protected. Next, we need to think about how we ensure that each component in the chassis can trust the other components and communicate with them in a secure manner.

As noted previously, each Management Module and iDRAC have a unique Identity certificate, signed by a Dell Certificate Authority. These certificates are a key part of the trust establishment between the various iDRAC and Management Modules inside the MX7000. Each system on startup will verify the installed certificates against the Dell EMC Root CA to assure they are valid. Certificates that are corrupted or invalid will be unable to establish machine to machine trust with other devices in the chassis.

Once assembled and powered on the MX7000 Management Modules and iDRACs will automatically start the discovery and machine to machine trust establishment process. All network communication inside the MX7000 is done over private IPv6 VLANs. The addresses in these VLAN’s are stateless and based on the router advertisement from the Enclosure controller. To locate other devices over the 2^128 IPv6 address space, individual devices use mDNS announcements to broadcast their presence. As the iDRAC and Management Modules discover each other they begin the process of establishing machine to machine trust.

A Management Module or iDRAC wishing to access resources on a discovered iDRAC or Management Module will need to prove it is an authentic Dell device to gain access. A ECIES (Elliptic Curve Integrated Encryption Scheme) using a ECDH (Elliptic Curve Diffie-Hellman) key exchange is used to pass the public portion of a “clients” certificate to the discovered “server”. The server will validate the certificate chain of the public certificate against the Dell root CA. If valid, the server will use the public key present in the certificate to form the shared symmetric key. The final piece of the

puzzle is giving the server a way to validate that the client actually has access to the private key for the unique identity certificate. To do this, the server uses a technique called “proof-of-possession” that is specified in RFC 7800. The proof-of-possession verification assures the server that the client has both public and private portions of a valid Dell Identity certificate. Having fully vetted the client, the server will provide the client with a temporary JWT (Java Web Token) that the server has signed and the client can use to access the resources of the server.

Encrypting management network traffic

In previous versions of PowerEdge servers, communications between devices with a chassis was expected to be secure due to the 'physical' security of a private internal network. A iDRAC blade would automatically become part of the chassis group by being physically inserted into the chassis. Communications to the sled from the CMC were done over telnet, HTTP, and other unencrypted channels. The authentication between the processes on the devices was often common shared passwords or other such preprogrammed credentials. While this was an extremely fast and robust design, Dell EMC has evaluated these processes and identified security concerns, concerns that have been addressed in the MX7000.

No longer are communications made over "clear text" HTTP connections, the new Redfish interface used in the chassis is done completely over HTTPS. The encryption of the REST information prevents packet snooping by other devices on the network. Previous multitenant chassis could be compromised by a malicious user, attempting to steal information from others on the same chassis. With the encrypted HTTPS communications this is no longer possible.

HTTPS is not the only communications path updated on the MX7000, all communications between iDRACs and Management Modules inside the chassis is encrypted. Linux sockets between iDRAC and the Management Module are encrypted using ECC (Elliptic Curve Cryptography). Communications over network sockets is possible only after iDRAC and Enclosure Controller have established bidirectional machine to machine trust. Only when both sides have vetted the other, are the connections established, using the keys transferred during trust establishment. This protects data passed between the devices, preventing snooping, as well as blocking attackers from pretending to be a Dell EMC device and accessing data.

Another issue addressed in the MX7000 is the use of a common default or fixed "hidden" user accounts with passwords programmed into the firmware. A fixed username/password know to the software allowed each device to quickly access and configure others without requiring the user interaction. The pitfalls of a common shared passwords are well documented and to avoid these issues the new MX7000 chassis uses unique, short duration and stateless token authentication. Unlike the normal username and password tokens are not tied to an actual user account on a device. In the MX7000 the iDRAC can issue an admin token to the MSM for reading/changing configuration without effecting user based authentication from its GUI. The MSM does not need a 'user' account and the new automated machine to machine trust assures the iDRAC is talking to an authentic MSM. Since the MSM is now a trusted administrator this presents all sorts of new possibilities.
Conclusion

Conclusions

Throughout this paper we have demonstrated how the PowerEdge MX solution has a robust security protocol and architecture. By implementing a secure boot process within the MX7000 we ensure that the system starts running only if the code passes integrity checks. Subsequently runtime security measures ensure that the system remains safe from malicious hacking attempts. And additionally, a validated network security measure ensures that everything in the chassis is a system running trusted code. These enhanced security measures use best-in-class tools to protect customer systems, and we believe that this new chassis represents the most secure chassis management system in the industry.