Mazama System Administration Infrastructure

Jonathan Sparks
Cray Inc

ABSTRACT: *Mazama is Cray's project to improve system administration for large scale MPP systems. It was developed under contract for the BlackWidow project and is being extended to Cray's other MPP systems.*

KEYWORDS: BlackWidow, Mazama, HSS, System Administration

1. Introduction

The Mazama system administration package is to be the common administration infrastructure for all future Cray MPP systems. At the time of Cray BlackWidow release partition Cray XT administrative functionality will be addressed by Mazama. Our strategy is to not disrupt current Cray XT administrative functionality in a revolutionary way. Functionality that will be common includes boot, alarm presentation and system dump tools.

2. Features

The functionality of Mazama is divided into two broad categories: system administration and system governance. System administration is used to describe day-to-day tasks including booting, image management, attributes and log management. System governance describes tasks that are longer term and includes configuration of hardware and services, provisioning of system resources, system failure prediction and analysis and file system management.

3. Architectural Representation

The System Management Workstation (SMW) functions as the system and maintenance workstation for the whole system, providing a single view of the configuration and central control. The SWM connects various system controllers together via a dedicated Ethernet network, known as the Hardware Supervisory System (HSS) network. Commands and system actions are routed to specific Controllers by means of the event router. The event router handles all the communication and routing of the command payload; system daemons have the task of performing the requested operations. The SMW also acts as the repository of system data, boot images and servers as the main authoritative controller of system tasks.

The administration architecture has been designed with the following objectives in mind:

- Facilitate fault analysis on large-scale systems.
- To allow system administration at scale, simplifying the per-node configuration as well as providing an aggregate view of configuration and control.
- To allow the administrator low-level access when required.
- To keep the high-speed network clear for application use, by using the HSS network for control functions.
- Delegate administration responsibility to the local controllers, leaving higher-level agents to define actions and policy.

For example, the boot process, Mazama will instruct each Cabinet Controller (CC) to boot a local set of nodes with an image (defined as OS + kernel+ parameters) and wait for the response. The Cabinet Controller is tasked with the mechanics of boot; initializing components, loading the boot image into memory via the Blade Controller (BC), taking the processor out of reset and presenting the node state back to Mazama.

Mazama is built on the concept of delegation of system administration and control to semi-autonomous agents such as the Cabinet Controller (CC), Blade Controller (BC) and Router Controller (RC) agent, as opposed to a centralized model, where a central authority is responsible for command and control.

The BlackWidow system contains a collection of semi-autonomous agents (sysd's) that communicate with one another to perform system control and monitoring. These agents communicate with one another via a message protocol over a dedicated Ethernet network, known as the Hardware Supervisory System (HSS).

The BlackWidow system employee's these agents to allow more flexibility for growth of capacity and function. In addition, the system can be more maintainable since each part can change so long as it continues to support its external interfaces.

Organization autonomy: Controlling agents want administrative control of facilities is critical to their effectiveness. To make this type of control scale most agents will require operational control of their resources and data. Increasingly, computer systems are being designed to reflect is organisational structure. There are several technological reasons for building a system administration infrastructure using delegation of actions to lower agents. It is sometimes not feasible to build a centralized system with the required capacity, response-time, or availability of a distributed system.

Typically these are:

- Capacity: Some applications have needs far in excess of resources allocated to a single service.
- Response time: Putting processing and the needed data closer to the requestor will increase response times.
- Locality: Storing local data and not disseminating it globally help in the containment of data.

Given the size of Cray MPP systems now and projected, a more decentralized approach has been taken with Mazama and HSS. Each Cabinet Controller acts as a semi-autonomous agent carrying out system administrative actions as directed by upper-level agents.

4. Features

This section describes the main features of the Mazama architecture that support system administration.

Command and Control

Command and Control gives the system administrator the ability to control hardware components (nodes, routers, blades). The administrator sends commands to the Cabinet Controller (CC), Blade Controller (BC) and/or Router Controller (RC). Depending on the type of command, if it's a complex operation such as boot, the CC may perform several independent operations to affect system boot then send a single response back to the controlling agent. This is an example of how Mazama delegate's command to sub-agents, Mazama does not need to know the internals of boot, just the need to perform the operation.

Monitoring

Monitoring means the ability to display and collect system information. This information is reported back to a central data store (Database) on the SMW, where various applications make use of the data.

For example:

- System state
- Network information
- Hardware and software errors
- System usage

Image Management

Image Management allows the system administrator to create, delete and upgrade OS distributions for deployment to compute nodes. The system administrator creates base OS distributions on the SMW, and applies site modifications to suite user/application demands. The resultant images create what is termed a bootable package, which is used for deployment to the nodes for booting. Another feature of Image Management will be the specialization of images to perform dedicated roles within the system, the addition and removal of system services with images will provide functionality pertinent to their role for example, adding Lustre will create an IO node.

Partition Management

Partition Management allows the administrator to assign compute nodes, OS distributions, kernels and configuration parameters to a single entity and used this entity for administration purposes such as startup, shutdown or status query. Partition attributes will be extended to include OS services and hardware characteristics as a means to fully define the context and usage of a partition.

Scheduling Attributes

Attributes are used to describe the characteristics of the smallest schedulable entities in the system. They provide a more sophisticated way to select the PEs on which to schedule jobs. Rather then requesting just the first *N* PEs that the system can allocate, a user can actually describe the characteristics of the PEs that they wish to run on.

There are two types of Mazama attributes; node attributes and socket attributes. Node attributes describe the attributes of all schedulable entities on a node. Within a node there are a number of sockets in which CPUs or FPGAs can be plugged into. Some attributes can be ascribed to individual sockets because they may vary form socket to socket. These are called socket attributes.

Reliability, Availability and Serviceability (RAS)

Mazama provides a system of software components that allow the administrator to:

- Monitor and report on hardware and software components.
- Monitor all logs and errors generated with the system.
- Filter all messages generated with the cabinet and based on site definable policies are sent bas to the SMW for storage and analysis.

Configuration

Configuration provides a system view of configuration, control and status. It includes a complete view of administration data within the system, in terms of a logical view (partitions, processors and images) and a physical view (represented by HSS).

5. Database

The Mazama database is used to keep track of system configuration and state. As shown in Figure 1, various external components can access this information via registration for administration events or by calling the Mazama database API

.

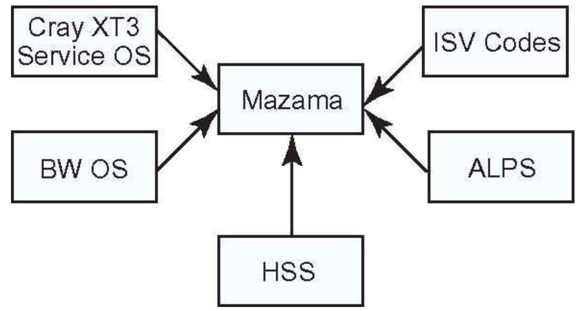

Figure 1. Mazama Database Interaction

The database represents both logical and physical views of the system; the logical view being centred on administration needs, and the physical around actual hardware components. ALPS (the Application Level Placement Scheduler) will query the database for node information and membership, independent software vendors (ISV) codes such as TotalView can look up node load, and Mazama will interact with HSS for command and control functionality.

From a high-level, the database reflects the notion of administration domains, each comprised of a single system instance. That means that one administration framework can administer multiple systems from a single point. Each system is further categorized into managed physical and logical entities; these entities define the system components, collections and attributes (for example, CPU, processors and state.)

HSS agents exist on each Cabinet Control (CC), Blade Controller (BC) and Router Controller (RC) and are responsible for in-cabinet control and monitoring component hardware. The administrative interface to this system is via the event mechanism and by database references. Each cabinet control processor maintains an independent database of components and state.

An aggregate system view of various HSS will also be available on the SMW. Changes in component state are communicated to Mazama by database table updates and events. It is the responsibility of the event router to propagate the notification event to all registered programs, one of which is Mazama. Events can be generated by Mazama, when a state change is administratively required (boot) or by HSS when a component state change is monitored (for example, voltage drop).

6. HSS

Hardware Supervisory system (HSS) will be common between Cray BlackWidow systems and future Cray systems. HSS is the backbone support network that ruins independent of the system fabric and controller machine hardware configuration and operation.

At a high level, the Hardware Supervisory System (HSS) can be thought of as a hardware management hierarchy with the System Management Workstation serving as the top-level manager. The major subordinate components of the HSS hierarchy include cabinet level controllers (CC), an intermediate entry for each blade, the Blade Host Controller either a Blade Controller (BC) or Router Controller (RC).

7. About the Author

Jonathan (Bill) Sparks is the manager of the system administration group developing the administration infrastructure for BlackWidow and beyond. He can be reached at jsparks@cray.com